

A Study of Muon Neutrino Disappearance Using the Fermilab Main Injector Neutrino Beam

P. Adamson,^{9,18} C. Andreopoulos,²³ K. E. Arms,¹⁹ R. Armstrong,¹² D. J. Auty,²⁷ S. Avvakumov,²⁶ D. S. Ayres,¹ B. Baller,⁹ B. Barish,⁵ P. D. Barnes Jr.,¹⁷ G. Barr,²¹ W. L. Barrett,³¹ E. Beall,^{1,19,*} B. R. Becker,¹⁹ A. Belias,²³ R. H. Bernstein,⁹ D. Bhattacharya,²² M. Bishai,⁴ A. Blake,⁶ B. Bock,²⁰ G. J. Bock,⁹ J. Boehm,¹⁰ D. J. Boehnlein,⁹ D. Bogert,⁹ P. M. Border,¹⁹ C. Bower,¹² E. Buckley-Geer,⁹ A. Cabrera,^{21,†} S. Cavanaugh,¹⁰ J. D. Chapman,⁶ D. Cherdack,³⁰ S. Childress,⁹ B. C. Choudhary,^{9,‡} J. H. Cobb,²¹ S. J. Coleman,³² A. J. Culling,⁶ J. K. de Jong,¹¹ M. Dierckxsens,⁴ M. V. Diwan,⁴ M. Dorman,^{18,23} D. Drakoulakos,² T. Durkin,²³ S. A. Dytman,²² A. R. Erwin,³³ C. O. Escobar,⁷ J. J. Evans,²¹ E. Falk Harris,²⁷ G. J. Feldman,¹⁰ T. H. Fields,¹ R. Ford,⁹ M. V. Frohne,^{3,§} H. R. Gallagher,³⁰ A. Godley,²⁵ J. Gogos,¹⁹ M. C. Goodman,¹ P. Gouffon,²⁴ R. Gran,²⁰ E. W. Grashorn,^{19,20} N. Grossman,⁹ K. Grzelak,²¹ A. Habig,²⁰ D. Harris,⁹ P. G. Harris,²⁷ J. Hartnell,²³ E. P. Hartouni,¹⁷ R. Hatcher,⁹ K. Heller,¹⁹ A. Himmel,⁵ A. Holin,¹⁸ C. Howcroft,⁵ J. Hysten,⁹ D. Indurthy,²⁹ G. M. Irwin,²⁶ M. Ishitsuka,¹² D. E. Jaffe,⁴ C. James,⁹ L. Jenner,¹⁸ D. Jensen,⁹ T. Kafka,³⁰ H. J. Kang,²⁶ S. M. S. Kasahara,¹⁹ M. S. Kim,²² G. Koizumi,⁹ S. Kopp,²⁹ M. Kordosky,^{32,18} D. J. Koskinen,¹⁸ S. K. Kotelnikov,¹⁶ A. Kreymer,⁹ S. Kumaratunga,¹⁹ K. Lang,²⁹ A. Lebedev,¹⁰ R. Lee,^{10,¶} J. Ling,²⁵ J. Liu,²⁹ P. J. Litchfield,¹⁹ R. P. Litchfield,²¹ L. Loiacono,²⁹ P. Lucas,⁹ W. A. Mann,³⁰ A. Marchionni,⁹ A. D. Marino,⁹ M. L. Marshak,¹⁹ J. S. Marshall,⁶ N. Mayer,^{12,20} A. M. McGowan,^{1,19,**} J. R. Meier,¹⁹ G. I. Merzon,¹⁶ M. D. Messier,¹² C. J. Metelko,²³ D. G. Michael,^{5,††} R. H. Milburn,³⁰ J. L. Miller,^{15,††} W. H. Miller,¹⁹ S. R. Mishra,²⁵ A. Mislivec,²⁰ C. D. Moore,⁹ J. Morfín,⁹ L. Mualem,^{5,19} S. Mufson,¹² S. Murgia,²⁶ J. Musser,¹² D. Naples,²² J. K. Nelson,³² H. B. Newman,⁵ R. J. Nichol,¹⁸ T. C. Nicholls,²³ J. P. Ochoa-Ricoux,⁵ W. P. Oliver,³⁰ T. Osiecki,²⁹ R. Ospanov,²⁹ J. Paley,¹² V. Paolone,²² A. Para,⁹ T. Patzak,⁸ Ž. Pavlović,²⁹ G. F. Pearce,²³ C. W. Peck,⁵ E. A. Peterson,¹⁹ D. A. Petyt,¹⁹ H. Ping,³³ R. Pittam,²¹ R. K. Plunkett,⁹ D. Rahman,¹⁹ R. A. Rameika,⁹ T. M. Rauffer,^{23,21} B. Rebel,⁹ J. Reichenbacher,¹ D. E. Reyna,¹ P. A. Rodrigues,²¹ C. Rosenfeld,²⁵ H. A. Rubin,¹¹ K. Ruddick,¹⁹ V. A. Ryabov,¹⁶ R. Saakyan,¹⁸ M. C. Sanchez,¹⁰ N. Saoulidou,⁹ J. Schneps,³⁰ P. Schreiner,³ V. K. Semenov,¹³ S.-M. Seun,¹⁰ P. Shanahan,⁹ W. Smart,⁹ V. Smirnitisky,¹⁴ C. Smith,^{18,27} A. Sousa,^{21,30} B. Speakman,¹⁹ P. Stamoulis,² M. Strait,¹⁹ P.A. Symes,²⁷ N. Tagg,^{30,21} R. L. Talaga,¹ E. Tetteh-Lartey,²⁸ J. Thomas,¹⁸ J. Thompson,^{22,††} M. A. Thomson,⁶ J. L. Thron,^{1,‡‡} G. Tinti,²¹ I. Trostin,¹⁴ V. A. Tsarev,¹⁶ G. Tzanakos,² J. Urheim,¹² P. Vahle,^{32,18} V. Verebryusov,¹⁴ B. Viren,⁴ C. P. Ward,⁶ D. R. Ward,⁶ M. Watabe,²⁸ A. Weber,^{21,23} R. C. Webb,²⁸ A. Wehmann,⁹ N. West,²¹ C. White,¹¹ S. G. Wojcicki,²⁶ D. M. Wright,¹⁷ T. Yang,²⁶ H. Zheng,⁵ M. Zois,² and R. Zwaska⁹

(The MINOS Collaboration)

¹Argonne National Laboratory, Argonne, Illinois 60439, USA

²Department of Physics, University of Athens, GR-15771 Athens, Greece

³Physics Department, Benedictine University, Lisle, Illinois 60532, USA

⁴Brookhaven National Laboratory, Upton, New York 11973, USA

⁵Lauritsen Laboratory, California Institute of Technology, Pasadena, California 91125, USA

⁶Cavendish Laboratory, University of Cambridge, Cambridge CB3 0HE, United Kingdom

⁷Universidade Estadual de Campinas, IF-UNICAMP, CP 6165, 13083-970, Campinas, SP, Brazil

⁸APC – Université Paris 7 Denis Diderot, Paris Cedex 13, France

⁹Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

¹⁰Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA

¹¹Physics Division, Illinois Institute of Technology, Chicago, Illinois 60616, USA

¹²Indiana University, Bloomington, Indiana 47405, USA

¹³Institute for High Energy Physics, Protvino, Moscow Region RU-140284, Russia

¹⁴High Energy Experimental Physics Department, Institute of Theoretical and Experimental Physics, B. Chermushkinskaya, 25, 117218 Moscow, Russia

¹⁵Physics Department, James Madison University, Harrisonburg, Virginia 22807, USA

¹⁶Nuclear Physics Department, Lebedev Physical Institute, Leninsky Prospect 53, 117924 Moscow, Russia

¹⁷Lawrence Livermore National Laboratory, Livermore, California 94550, USA

¹⁸Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, United Kingdom

¹⁹University of Minnesota, Minneapolis, Minnesota 55455, USA

²⁰Department of Physics, University of Minnesota – Duluth, Duluth, Minnesota 55812, USA

²¹Subdepartment of Particle Physics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom

²²Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, Pennsylvania 15260, USA

²³Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire, OX11 0QX, United Kingdom

²⁴Instituto de Física, Universidade de São Paulo, CP 66318, 05315-970, São Paulo, SP, Brazil

²⁵Department of Physics and Astronomy, University of South Carolina, Columbia, South Carolina 29208, USA

²⁶Department of Physics, Stanford University, Stanford, California 94305, USA

²⁷Department of Physics and Astronomy, University of Sussex, Falmer, Brighton BN1 9QH, United Kingdom

²⁸Physics Department, Texas A&M University, College Station, Texas 77843, USA

²⁹Department of Physics, University of Texas, Austin, Texas 78712, USA

³⁰Physics Department, Tufts University, Medford, Massachusetts 02155, USA

³¹Physics Department, Western Washington University, Bellingham, Washington 98225, USA

³²Department of Physics, College of William & Mary, Williamsburg, Virginia 23187, USA

³³Physics Department, University of Wisconsin, Madison, Wisconsin 53706, USA

(Dated: May 23, 2018)

We report the results of a search for ν_μ disappearance by the Main Injector Neutrino Oscillation Search [1]. The experiment uses two detectors separated by 734 km to observe a beam of neutrinos created by the Neutrinos at the Main Injector facility at Fermi National Accelerator Laboratory. The data were collected in the first 282 days of beam operations and correspond to an exposure of 1.27×10^{20} protons on target. Based on measurements in the Near Detector, in the absence of neutrino oscillations we expected 336 ± 14 ν_μ charged-current interactions at the Far Detector but observed 215. This deficit of events corresponds to a significance of 5.2 standard deviations. The deficit is energy dependent and is consistent with two-flavor neutrino oscillations according to $|\Delta m^2| = 2.74_{-0.26}^{+0.44} \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta > 0.87$ at 68% confidence level.

I. INTRODUCTION

There is now strong evidence that ν_e produced in the Sun [2, 3] and $\bar{\nu}_e$ produced in nuclear reactors [4] change flavor. There is additional and compelling evidence that $\nu_\mu/\bar{\nu}_\mu$ produced in the atmosphere [5, 6, 7] and more recently by accelerators [1, 8] disappear while propagating through both the atmosphere and the Earth over distances of 250-13000 km. This behavior is consistent with three-flavor neutrino oscillations induced by mixing between non-degenerate mass eigenstates ν_1, ν_2, ν_3 and the states ν_e, ν_μ, ν_τ created in weak interactions. The flavor and mass eigenstates are related by a unitary matrix U_{PMNS} [9, 10] which is typically expressed in terms of three angles $\theta_{12}, \theta_{23}, \theta_{13}$ and a CP-violating phase δ . The energy and distance scales at which oscillations occur are determined by the difference in squared masses, $\Delta m_{ij}^2 = m_i^2 - m_j^2$, of the ν_1, ν_2, ν_3 .

The data suggest that the phenomenology of solar and reactor neutrinos is driven by a squared-mass splitting $7.6 \leq \Delta m_\odot^2 \leq 8.6 \times 10^{-5} \text{ eV}^2/c^4$ (68% confidence level – C.L.) [2] whereas the behavior of atmospheric and accelerator produced ν_μ is determined by a much larger splitting $1.9 \leq |\Delta m_{\text{atm}}^2| \leq 3.0 \times 10^{-3} \text{ eV}^2/c^4$ (90% C.L.) [5]. In this case muon-neutrino disappearance may be described as two-flavor neutrino oscillations:

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2 2\theta \sin^2 \left(1.27 \frac{\Delta m^2 [\text{eV}^2] L [\text{km}]}{E [\text{GeV}]} \right) \quad (1)$$

where Δm^2 is the effective difference of squared masses, θ is a mixing angle between the mass and weak eigenstates, E is the neutrino energy and L is the distance from the neutrino production point to the observation point. The analysis reported here is conducted in the two-neutrino framework, and we can tentatively identify $\Delta m^2 \approx \Delta m_{\text{atm}}^2 \approx \Delta m_{32}^2$ and $\sin^2 2\theta = \sin^2(2\theta_{23})$ in the limit $P(\nu_\mu \rightarrow \nu_e) = 0$. More precisely, the experiment measures $\Delta m^2 = \sin^2 \theta_{12} \Delta m_{31}^2 + \cos^2 \theta_{12} \Delta m_{32}^2$ where effects of order $\sin^2(\theta_{13})$ have been neglected [11].

The Main Injector Neutrino Oscillation Search (MINOS) is a long-baseline, two-detector neutrino oscillation experiment that uses a muon-neutrino beam produced by the Neutrinos at the Main Injector [12, 13] (NuMI) facility at Fermi National Accelerator Laboratory (FNAL). NuMI is able to provide a range of neutrino energies but has

*Now at Cleveland Clinic, Cleveland, Ohio 44195, USA.

†Now at APC – Université Paris 7 Denis Diderot, Paris Cedex 13, France.

‡Now at Dept. of Physics and Astrophysics, University of Delhi, Delhi 110007, India.

§Now at Holy Cross College, Notre Dame, Indiana 46556, USA.

¶Now at Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, Massachusetts 02420, USA.

**Now at Physics Department, St. John Fisher College, Rochester, New York 14618 USA.

††Deceased.

‡‡Now at Nuclear Nonproliferation Division, Threat Reduction Directorate, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA.

mostly been operated in the “low-energy” (LE10/185kA of Tab. I) configuration, which maximizes the neutrino flux at $E \approx 3$ GeV. Neutrinos are observed by two functionally identical detectors, located at two sites, the Near Detector (ND) at FNAL and the Far Detector (FD) in the Soudan Underground Laboratory in Minnesota. The detectors are separated by 734 km and are designed to detect muon-neutrino charged-current interactions for $1 < E < 100$ GeV. The characteristic $L/E \approx 245$ km/GeV allows MINOS to rigorously test the oscillation hypothesis and make precision measurements of the Δm^2 and $\sin^2 2\theta$ mixing parameters governing muon-neutrino disappearance at the atmospheric neutrino mass-scale.

MINOS employs two detectors to significantly reduce the effect that systematic uncertainties associated with the neutrino flux, cross-sections and detector efficiency have upon the ν_μ disappearance measurement. Data collected by the Near Detector in several different configurations of the NuMI beam (Tab. I and Fig. 2) were used to constrain the simulation of the neutrino flux and the detector response to neutrino interactions. Four independent analyses were used to predict the energy spectrum expected at the Far Detector for the case that ν_μ do not disappear. The analyses were fully developed using Near Detector data only. After we were satisfied with the procedure, the energy spectrum measured at the Far Detector was inspected and then compared to the predictions; oscillations cause an energy dependent suppression of the ν_μ charged-current rate. A fit to the Far Detector data was performed by incorporating Eq. 1 into the Far Detector prediction. The fit also included the most significant sources of systematic uncertainty and resulted in a measurement of the oscillation parameters.

MINOS began collecting NuMI beam data in March 2005. The Far Detector dataset used in the analysis reported here was recorded between May 20, 2005 and February 25, 2006. The total exposure is 1.27×10^{20} protons on target (POT) and only includes data collected while NuMI was operating in the low-energy beam configuration. This exposure allows MINOS to measure Δm^2 and $\sin^2 2\theta$ with a precision that is comparable to the best existing measurements [5, 8]. The results reported here use the same dataset and are identical to those reported in [1] but include a more detailed description of the experiment, analysis and results. Since the publication of [1], MINOS has accumulated a total of 3.5×10^{20} POT through July 2007 and preliminary results based on 2.5×10^{20} POT have been presented [14]. Analysis of the full 3.5×10^{20} POT dataset is ongoing, and accumulation of further data is foreseen.

We begin in Sec. II by describing the neutrino beam-line design, operation and simulation. The neutrino detectors are described in Sec. III along with the calibration procedure and the simulation of neutrino interactions. We also discuss neutrino data collection and event reconstruction in Sec. III and the selection of ν_μ charged-current events in Sec. IV. The manner in which Near Detector neutrino data was used to constrain the neutrino flux calculation is presented in Section V. Four methods for predicting the Far Detector ν_μ charged-current spectrum in the absence of oscillations are discussed in Sec. VI. We conclude in Sec. VII with a description of the oscillation analysis and measurement of the parameters Δm^2 and $\sin^2 2\theta$.

II. THE NEUTRINO BEAM

The NuMI Neutrino Beam [12, 13] generates neutrinos mainly from the decays of pions and kaons produced in the NuMI target, with a smaller contribution from muon decays. A schematic diagram of the NuMI beam-line is shown in Fig. 1. Protons of 120 GeV/ c momentum are extracted from the Main Injector (MI) accelerator in a $10 \mu\text{s}$ spill, bent downward by 58 mrad to point at the Far Detector, and impinged upon the NuMI hadron production target. Positioned downstream of the target, two toroidal magnets called “horns” sign-select and focus the secondary mesons from the target, as shown in Fig. 2. The mesons are directed into a 675 m long evacuated volume, where they may decay to muons and neutrinos. At the end of the decay volume, a beam absorber stops the remnant hadrons. The absorber is followed by approximately 240 m of un-excavated rock which stops the tertiary muons, leaving only neutrinos. Figure 3 shows ν_μ charged-current spectra from three configurations of the NuMI beam. The following section provides a brief overview of the beam-line and instrumentation.

A. Beam Hardware and Performance

1. Primary Proton Beam

The Main Injector accepts batches of protons from the 8 GeV/ c Booster accelerator, which are then accelerated to 120 GeV/ c . In most MI cycles seven batches were accelerated. Protons were removed from the MI ring in two increments, with a pair of batches being sent to the Antiproton Accumulator for the Tevatron program, while the remaining five batches were directed into the NuMI primary-proton line and transported 350 m to the NuMI target. The NuMI extractions typically contained a total of 2.1×10^{13} protons with a cycle time of 2.2 – 2.4 s. Extractions of up to six batches and 2.5×10^{13} protons were achieved during the first year of operations. The proton beam centroid




FIG. 1: Plan and elevation views of the NuMI beam facility. A proton beam is directed onto a target, where the secondary pions and kaons are focused into an evacuated decay volume via magnetic horns. Ionization chambers at the end of the beam-line measure the secondary hadron beam and tertiary muon beam.




FIG. 2: In the NuMI neutrino beam, horns 1 and 2 are separated by 10 m. A collimating baffle upstream of the target protects the horns from direct exposure to misdirected proton beam pulses. The target and baffle system can be moved further upstream of the horns to produce higher energy neutrino beams [16]. The vertical scale is four times that of the horizontal (beam axis) scale.

at the target was stable to within ± 0.1 mm and the area of the beam-spot varied within the range $3.3 - 4.5$ mm². Further details of the primary proton beam delivery system and its performance are given in Ref. [15].

2. Target and Horns

The production target is a rectangular graphite rod, segmented longitudinally into 47 segments, or fins. The target dimensions are 6.4 mm in width, 15 mm in height and 940 mm in length. The beam spot size at the target is 1.2-




FIG. 3: Calculated rate of ν_μ charged-current interactions in the MINOS Near Detector, located 1040m from the NuMI target. The figures were made by combining the flux calculation discussed in II B with the cross-sections discussed in Sec. III E. Three spectra are shown, corresponding to the LE10/185kA, LE100/200kA, and LE250/200kA configurations of Tab. I.

1.5 mm [15]. A collimating baffle upstream of the target provides protection for the target, its cooling lines, as well as downstream beam components. The baffle is a 1.5 m long graphite rod with an 11 mm diameter inner bore.

The particles produced in the target are focused by two magnetic horns [17], shown schematically in Fig. 2. The 200 kA pulsed current produces a maximum 30 kG toroidal field which selects particles produced at the target by charge-sign and momentum. Measurements of the horns show the expected $1/r$ fall-off in the magnetic field to within a percent and the field at distances from the beam axis smaller than the inner conductors is observed to be less than 20 G. The absolute value of the current flowing through the horns was calibrated to within $\pm 0.5\%$ and was observed to vary less than 0.2% over the course of the data collection period. The horns have parabolic-shaped inner conductors [17], for which the focal length for point-to-parallel focusing of particles from the target is proportional to momentum [18, 19]. The alignment of the target and horn system relative to the beam axis was checked using the proton beam itself [20].

The relative longitudinal positions of the two horns and the target optimizes the momentum focus for pions and therefore the typical neutrino energy. To fine-tune the beam energy, the target is mounted on a rail-drive system with 2.5 m of longitudinal travel, permitting remote change of the beam energy without directly accessing the horns and target [16]. In its furthest downstream location, the target is cantilevered approximately 65 cm into the first parabolic horn. Characteristics of the predicted neutrino spectra from several target position and horn current settings are listed

Beam Configuration	Target Position (cm)	Horn Current (kA)	Peak $E_\nu \pm \text{r.m.s.}$ (GeV)	Exposure 10^{18} POT
LE10/0kA	10	0	7.4 ± 4.1^a	2.69
LE10/170kA	10	170	3.1 ± 1.1	1.34
LE10/185kA	10	185	3.3 ± 1.1	127.
LE10/200kA	10	200	3.5 ± 1.1	1.26
LE100/200kA	100	200	5.6 ± 1.5	1.11
LE250/200kA	250	200	8.6 ± 2.7	1.55

^aThe 0 kA “horn-off” beam is unfocused and has a broad energy distribution

TABLE I: Beam configurations and data sets used in this publication. The target position refers to the distance the target was displaced upstream of its default position inside the first focusing horn. The peak (i.e., most probable) neutrino energy E_ν is determined after multiplying the muon-neutrino flux predicted by the beam Monte Carlo simulation by charged-current cross-section. The r.m.s. refers to the root mean square of the peak of the neutrino energy distribution.

in Tab. I and ν_μ charged-current energy distributions are shown in Fig. 3. Data from the LE10/185kA configuration were used in the oscillation search. Data from the other beam configurations were used to constrain the beam Monte Carlo (see Sec. V).

3. Decay Volume and Absorber

Particles are focused by the horns into a 675 m long, 2 m diameter steel pipe, evacuated to 0.5 Torr to reduce meson absorption and scattering. This length is approximately the decay length of a 10 GeV pion. The entrance to the decay pipe is sealed by a two-piece aluminum-steel window. The central (radius < 50 cm) portion of the window is made of 1 mm thick aluminum and is strengthened by an outer (radius > 50 cm) section made of 1.8 cm thick steel. The design reduces scattering in the window while maintaining vacuum integrity. The decay volume is surrounded by 2.5-3.5 m of concrete shielding. At the end of the decay volume is a beam absorber consisting of a water-cooled aluminum and steel core followed by steel and concrete blocks.

4. Instrumentation

The primary proton beam position is monitored along the transport line by 24 capacitive beam position monitors (BPMs), the beam intensity by two toroidal beam current transformers and 44 loss monitors, and the beam position and spot size by ten retractable segmented foil secondary emission monitors (SEMs) [21]. During normal operations only the last SEM upstream of the target is inserted in the beam. The absolute toroid uncertainty was determined to be $\pm 1.0\%$ by precision current pulses and monitored for drift throughout the run by comparison of the toroids in the NuMI line with monitors in the Main Injector. The beam position and size at the target are measured to within $\pm 50 \mu\text{m}$ [15].

In the target hall, the current flowing in the horns was monitored spill-to-spill, as was the temperature of the current-delivering strip-line. Both were seen to change by approximately 0.2% due to thermal variations in the target hall. The temperature of the upstream collimating baffle was continuously monitored, as this was observed to be a good measure of the proton beam halo scraping on the baffle. During NuMI operations it was found that approximately 0.3% of the beam was obstructed by the baffle.

Ionization chambers are used to monitor the secondary and tertiary particle beams [22]. An array is located immediately upstream of the beam absorber to monitor the remnant hadrons at the end of the decay pipe. There are also three muon monitoring stations, one downstream of the absorber, one after 12 m of rock, and a third after an additional 18 m of rock. The muon stations confirmed the relative alignment of the neutrino beam direction to approximately $20 \mu\text{rad}$ throughout the run as shown in Fig. 4. This is expected to limit spectral variations at the Far Detector to less than 0.2%. In addition, the charge (per spill per proton on target) deposited in the muon monitors varied by only $\pm 2\%$, indicating a similar level of stability in the neutrino flux.

The relative position of a point on the surface above the Far Detector and the NuMI target were measured by GPS survey with an accuracy of 1 cm in the horizontal plane, and 6 cm in the vertical [12, 23]. Translating the surface coordinates to the underground Far Detector hall introduces an additional uncertainty of 70 cm in the vertical. The

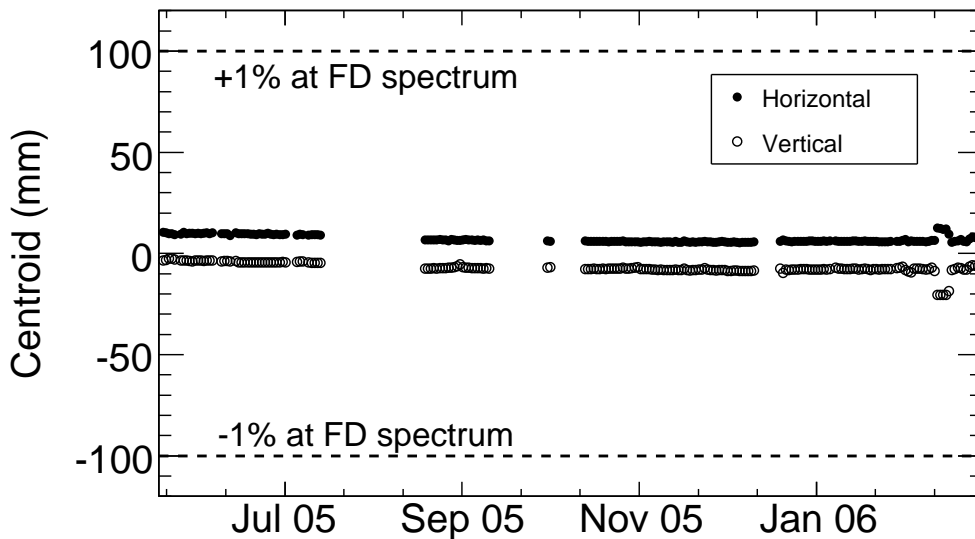


FIG. 4: Horizontal and vertical beam centroid measured by the first muon station. The observed variations in the centroid correspond to $20 \mu\text{rad}$ variations in the beam angle. Such deviations are expected to have less than 0.2% effect on the ν_μ energy spectrum at the Far Detector .

positions of beam-line components with respect to the target were surveyed and are known with an accuracy of 0.5 mm (95% C.L.). These data indicate that the beam angle with respect to the ideal Near to Far Detector vector is less than $15 \mu\text{rad}$ (68% C.L.). This introduces a negligible uncertainty into the calculation of the neutrino flux at the Far Detector.

B. Simulation of the Neutrino Beam


The neutrino beam-line is modeled in three stages: (1) a simulation of the hadrons produced by 120 GeV/c protons incident on the NuMI target, (2) the propagation of those hadrons and their progeny through the magnetic focusing elements, along the 675 m decay pipe and into the primary beam absorber allowing for decay of unstable particles and (3) the calculation of the probability that any produced neutrino traverses the Near and Far detectors.

The production of secondary mesons in the NuMI target was calculated using the FLUKA05 [24] Monte Carlo. Particles exiting the target are recorded and later propagated in a GEANT3 [25] simulation of the NuMI beam-line. The simulation describes the magnetic focusing horns, surrounding shielding and decay pipe. The magnetic field inside the horns was modeled as a perfect $1/r$ toroidal field. The magnetic field in the inner conductor of the horn was calculated assuming that the current was uniformly distributed throughout the 3 mm thickness of the aluminum conductor. This assumption was motivated by the large skin depth $\delta \approx 7.7$ mm for the approximately 1 ms current pulse. The 10-30 G magnetic field in the inner apertures of the horns was neglected, as was the approximately 1 G field in the decay pipe due to residual magnetization of the iron vessel. The GEANT-FLUKA code is used to describe hadronic interactions in the beam-line and the associated production of secondary particles as well as the full particle decay chains. Decays in which a neutrino is produced are saved and later used as input for neutrino event simulation in the Near and Far Detectors.

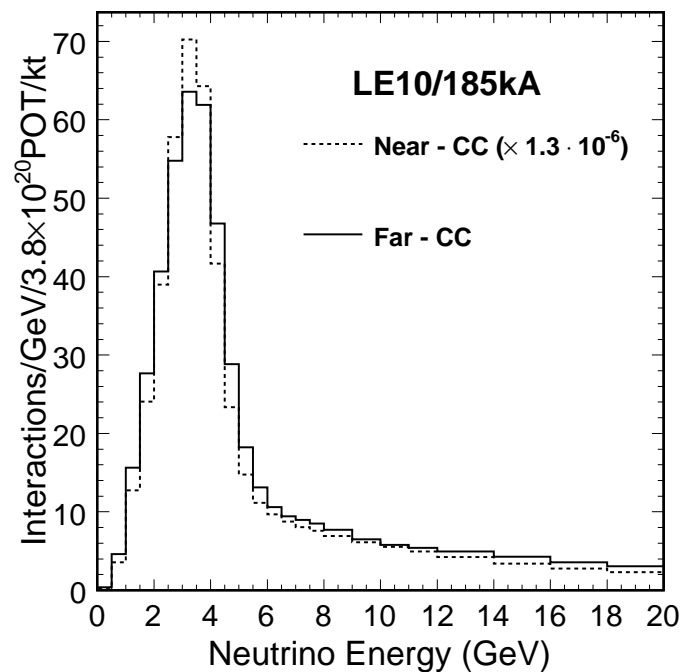
The neutrino event simulation uses each decay recorded by the NuMI simulation with a probability and neutrino energy determined by the decay kinematics and the (randomly chosen) trajectory through the Near or Far Detector. For two-body $\pi/K \rightarrow \mu\nu$ decays of relativistic mesons, the neutrino energy is given by

$$E_\nu \approx \frac{(1 - \frac{m_\mu^2}{M^2})E}{1 + \gamma^2 \tan^2 \theta_\nu} \quad . \quad (2)$$

where m_μ and M are the muon and parent hadron masses, E the parent hadron energy, $\gamma = E/M$ is the parent's Lorentz boost, and θ_ν is the angle (in the lab) between the neutrino and parent hadron directions. The neutrino is



(a)



(b)

FIG. 5: (a) Expected energy distribution of ν_μ charged-current interactions in the Near Detector. The contributions from π and K decays are shown separately. (b) Comparison of the neutrino energy spectrum at the Near and the Far Detectors. The two are not identical, due to solid angle differences between the two detectors (see Fig. 6).




FIG. 6: The differences in the Near Detector and Far Detector spectra are due to the different angular acceptance of the two detectors and the proximity of the Near Detector to the beam-line

forced to pass through either the Near or Far MINOS Detector, with probability for the particular meson decay given by

$$\frac{dP}{d\Omega_\nu} \approx \frac{1}{4\pi} \frac{4\gamma^2(1 + \tan^2 \theta_\nu)^{3/2}}{(1 + \gamma^2 \tan^2 \theta_\nu)^2}, \quad (3)$$

In deriving both expressions, $\beta = \sqrt{1 - (1/\gamma^2)} \approx 1$ is assumed. We have also accounted for the effect of μ polarization (e.g., induced in $\pi \rightarrow \mu\nu$) on the neutrino flux from μ decays. Three-body kaon decays are also included, but contribute $< 0.1\%$ to the ν_μ event rate.

C. Expected Neutrino Energy Spectra


Figure 3 shows the energy spectra of ν_μ charged-current interactions in the MINOS Near Detector in three of the beam configurations of Tab. I. In the LE10/185kA beam configuration, the neutrino in 87% of ν_μ charged-current interactions was produced by $\pi^+ \rightarrow \mu^+ \nu_\mu$, with K^+ decays contributing the additional 13%. K^0 and μ decays contribute $< 0.1\%$ to ν_μ event rate, though they do contribute significantly to the ν_e rate. The contributions of π and K to the neutrino energy spectrum is shown in Fig. 5. The predicted spectrum for the LE10/185kA beam is composed of 92.9% ν_μ , 5.8% $\bar{\nu}_\mu$ and 1.3% $\nu_e + \bar{\nu}_e$.

The uncertainties on the neutrino flux due to focusing effects are expected to be largest at the edges of apertures of the horns. The uncertainties are shown in Fig 7a-7b and described in more detail in Sec IID. Uncertainties at higher neutrino energies are smaller because higher energy neutrinos arise from high energy π and K mesons which pass through the field-free apertures of the horns without focusing.

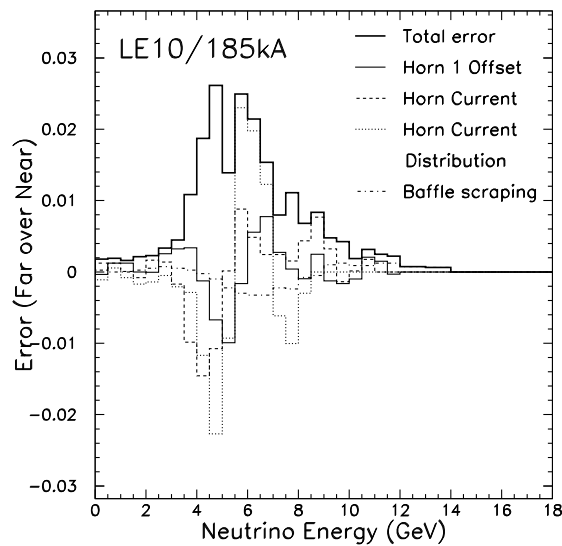
The neutrino energy spectra at the MINOS Near and Far Detectors are not identical, due to the different solid angles subtended by the two detectors and due to the energy-dependence of the meson decay locations. As indicated in Fig. 6, higher-energy pions decay further down the length of the decay volume, and hence a variety of decay angles θ_ν result in neutrinos which strike the Near Detector. The Far Detector, by contrast, can only be struck for very small-angle decays. By Eq. 2, the different allowed decay angles, as well as the close proximity of fast pion decays to the Near Detector, results in different energy spectra at the two detectors. These differences are conveniently characterized in terms of the Far/Near spectral ratio shown in Fig. 29. Thus, while the Near Detector in principle measures the neutrino energy spectrum in the absence of oscillations, corrections must be applied to this spectrum to derive a predicted spectrum at the Far Detector, as is discussed in Sec. VI.

D. Uncertainties In The Neutrino Flux

The dominant contribution to the uncertainty in the flux at both detector sites is caused by uncertainty in the yield of hadrons off the target as a function of p_z and p_T , the components of the hadron momentum along and transverse to the beam-line. The magnitude of the uncertainty on the energy spectra is difficult to estimate given the lack of hadron production data. As described in Sec. V, the data in the Near Detector can be used to constrain the beam flux calculation, and in particular the yield of hadrons off the target. The constrained calculation improves agreement with the Near Detector data and reduces uncertainties in the prediction of the Far Detector flux. The procedure used



(a)




(b)


FIG. 7: The effect that beam focusing uncertainties have on (a) the Near Detector ν_μ charged-current spectrum and (b) on the ratio (Far/Near) of ν_μ charged-current spectra at the two detectors. The lines show the fractional change in the number of events (a) or event ratio (b) expected in each energy bin due to a one standard deviation shift in various beam parameters. These include the horn and target alignment, proton beam halo scraping on the collimator baffle, knowledge of the horn current, and the modeling of the current distribution in the horn conductor. The total error is the sum in quadrature of the individual uncertainties.

to apply this constraint is described in Sec. V. Unless otherwise noted the Monte Carlo simulation will use the flux calculation constrained by the Near Detector data.

Systematic uncertainties on the predicted neutrino flux from beam focusing effects are readily calculable, and are constrained using data from the instrumentation in the primary, secondary, and tertiary beams. Figure 7a shows the expected uncertainty on the neutrino energy spectrum at the Near Detector. Focusing uncertainties produce spectral distortions in both detectors that occur predominantly in the portions of the neutrino energy spectrum which



(a)



(b)

FIG. 8: The MINOS Near (a) and Far (b) Detectors. The Far Detector consists of two functionally identical modules, only one of which is shown in the figure.

correspond to parent hadrons which cross the edges of the horn apertures (see Fig. 2). The largest beam focusing uncertainties include our knowledge of the absolute current flowing through the horns, the alignment of the horns and target to the rest of the beam-line [20], the shielding geometry, the uncertainty in the fraction of the proton beam scraping on the upstream collimating baffle (which acts as a target further upstream of the horns), and the modeling of the current distribution in the horns' inner conductors (e.g., skin-depth effect). The expected distortion in the ratio of the Far and Near Detector spectra are smaller (see Fig. 7b) than the individual Near or Far Detector uncertainties (Fig. 7a). Hence, use of the Near Detector to measure the spectrum results in a prediction for the spectrum at the Far Detector with smaller uncertainty.

III. THE NEUTRINO DETECTORS

MINOS consists of two neutrino detectors (shown in Fig. 8) separated by a long baseline. The Near Detector resides on the Fermilab site, 104m underground and 1040 m downstream of the NuMI target. The Far Detector is




FIG. 9: Optical readout of the MINOS detectors. Scintillation light is captured by wavelength-shifting fibers embedded in the scintillator and then transferred to multi-anode photomultipliers via clear optical fibers. The Far Detector uses 16-anode phototubes and the Near Detector uses 64-anode phototubes. In the Far Detector eight fibers are optically summed on a single pixel in a manner which permits the eight-fold ambiguity to be resolved during the event reconstruction.

located in the Soudan Underground Laboratory at a distance of 735 km from the NuMI target and 705 m beneath the surface. The detectors are capable of observing muon-neutrino and electron-neutrino charged-current and neutral-current interactions having a visible energy larger than about 500 MeV. Each detector has a toroidal magnetic field which is used to measure muon momenta via curvature and discriminate between ν_μ and $\bar{\nu}_\mu$. The two detectors were designed to respond to neutrino interactions in the same way so as to reduce the impact that uncertainties in the neutrino flux, cross-sections and detector acceptance have on oscillation measurements. In this section we will first describe the common features of both detectors, then those specific to the Far and Near Detectors, followed by a description of the data collection, event reconstruction, calibration and simulation. A more comprehensive description of the MINOS detectors is presented in [26].

A. Detector Technology

The MINOS detectors are magnetized steel/scintillator tracking/sampling calorimeters designed to measure muon-neutrinos produced by the NuMI beam. The active medium comprises 4.1 cm-wide, 1.0 cm-thick plastic scintillator strips arranged side and encased within aluminum sheets to form light-tight modules of 20 or 28 strips. Modules are combined to form scintillator planes which are then mounted on a 2.54 cm-thick steel absorber plate. The detectors are composed of a series of these steel/scintillator planes hung vertically at a 5.94 cm pitch. The scintillator strips in successive planes are rotated by 90° to measure the three dimensional event topology. The strips were produced in an extrusion process during which a 2.0 mm-wide, 2.0 mm-deep groove was driven along each strip. The entire strip, apart from the groove, was co-extruded with a 0.25 mm-thick TiO_2 doped polystyrene reflective layer. Scintillation light is

collected by a 1.2 mm-diameter wavelength-shifting fiber that was glued into the groove with optical epoxy. The fiber transports the light to an optical connector located on the edge of the scintillator layer where it is transferred to a clear polystyrene fiber and routed to multi-anode photomultipliers, as shown in Fig. 9. Light emitting diodes are used to produce controlled pulses of light during the data-taking to track gain changes in and characterize the response of the photomultipliers and electronics [27]. The data acquisition (DAQ) [28] and timing system synchronize and continuously read out the front-end electronics. Software triggering in the DAQ provides flexible event selection and data processing. GPS timestamps allow data from the two detectors to be synchronized with the beam pulses. The two detectors have different front-end electronics due to the disparate rates of neutrino interactions and cosmic-ray crossings at the two sites.

B. The Far Detector

The Far Detector has a mass of 5400 metric tons and consists of 486 steel plates arranged in two modules separated by a 1.1 m gap. The plates are 2.54 cm thick, 8 m wide octagons. A 15.2 kA magnet coil is routed through a hole in the center of each plane and induces an average field of 1.27 T in the steel. Each scintillator plane covers the 8 m octagon and consists of 192 scintillator strips read out from both ends by Hamamatsu 16-anode photomultipliers. The signals from eight strip-ends, separated by approximately 1 m within a single plane, are optically summed onto a single photomultiplier channel. The optical summing pattern is different for the two sides of the detector, allowing for resolution of the eight-fold ambiguity associated with each hit.

The front-end electronics [29] digitize signals from the photomultipliers with a 14 bit ADC (2 fC precision) when the common dynode signal exceeds approximately 0.25 photoelectrons and time-stamps them with a 1.5625 ns least significant bit. To reduce the electronics dead-time, hits are only digitized if more than one of the photomultipliers serving a contiguous group of 20 or 24 planes on one side of the detector are above threshold in coincidence. Data selected for further processing are pedestal suppressed and sent to a trigger farm where software triggering and data processing are performed and GPS timestamps are applied.

The DAQ has data buffering that is large enough to allow it to wait for the GPS time-stamp of the spill to arrive from the Near Detector over the Internet. This is used to form a bias-free beam trigger by recording all hits in the detector within a 100 μ s window around the time of the spill. To avoid splitting events, the window size is extended for each spill to ensure that the entire window is bounded on both sides by an activity free period of at least 156 ns. Finally, all hits in the 30 μ s prior to each trigger are added to the event to provide a mapping of channels that were engaged in digitization at the time of the trigger. Fake spill triggers are also generated to monitor backgrounds. When spill information is not available at the Far Detector site, and for all out-of-spill triggering, candidate events are formed from time sequential blocks of hits bounded on each side by at least 156 ns in which no detector activity occurred; trigger algorithms based on the spatial and energy clustering of hits in the detector are then applied to select events of potential interest. The integrated trigger rate in the Far Detector is typically 4 Hz and is dominated by cosmic-rays and single photoelectron noise.

C. The Near Detector

The 980 metric ton Near Detector consists of 282 steel/scintillator planes arranged in a single magnetized module. Each steel plate has a “squashed-octagon” shape 6.2 m wide and 3.8 m high with a 30×30 cm² hole offset 56 cm from the horizontal center of the detector to accommodate a magnet coil. The 40 kA carried by the coil induces a 1.17 T field at the neutrino beam center, located 1.49 m to the left of the coil. The coil and detector geometry was designed to provide a magnetic field in the region around the beam center that is similar to the field in the Far Detector. The Near Detector has two different types of scintillator planes: partially instrumented and fully instrumented. The planes are smaller than those in the Far Detector and are read out on only one side using Hamamatsu 64 anode photomultipliers. Fully instrumented Near Detector planes have 96 strips and cover a 13.2 m² area. These planes are attached to one in every five steel plates along the entire length of the detector. Partially instrumented planes are attached to four out of five plates in the upstream portion of the detector. These planes are comprised of 64 strips covering a 6.0 m² area.

The Near Detector is organized in two sections. The upstream 121 steel plates form the calorimeter, which is used to define the interaction vertex, find the upstream portion of muon tracks, and measure the energy of the neutrino induced hadronic shower. In the calorimeter each plate is instrumented with scintillator and the signals from each strip are read out independently. The downstream 161 plates are used as a muon spectrometer and only one in every five is instrumented. Furthermore, in the spectrometer, the signals from four strips are summed onto one electronics

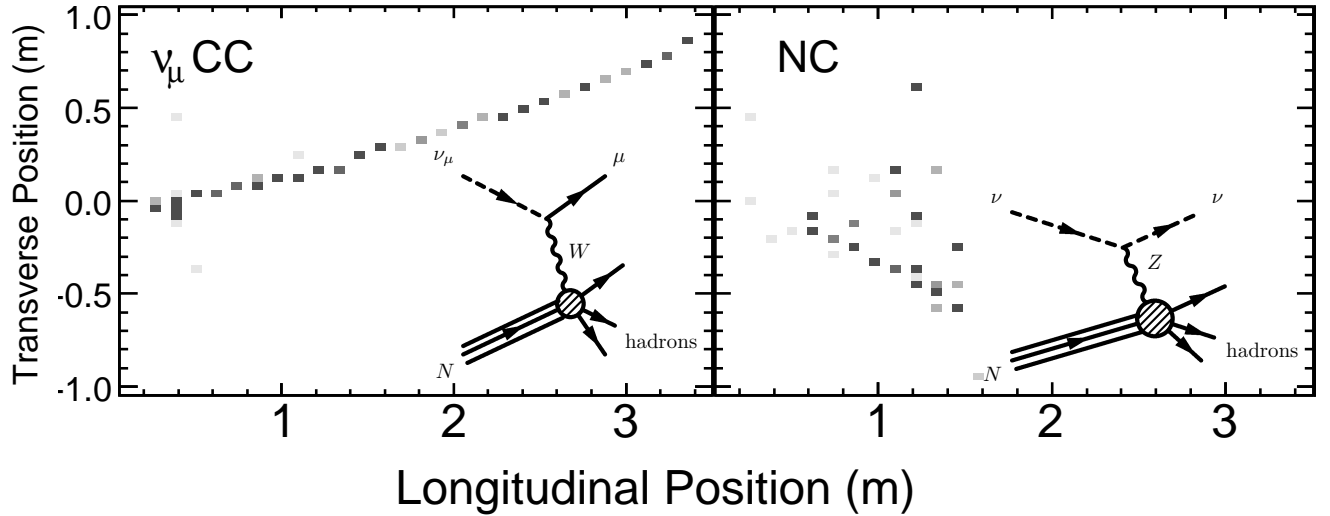


FIG. 10: Muon-neutrino charged-current (ν_μ charged-current) and neutral-current (NC) events simulated in the MINOS detector. Shaded rectangles indicate energy depositions (hits) in the detector’s strips. The strips are organized in planes transverse to the beam axis and the strips in successive planes are rotated by $\pm 90^\circ$ to provide two orthogonal views of the event. This figure shows the data from one view.

channel. The four-fold ambiguity is resolved in the event reconstruction program by extrapolating the muon track found in the upstream portion of the detector.

At the typical beam intensity of 2.2×10^{13} POT/spill, an average of 16 (44) neutrino interactions¹ occurred in the Near Detector during each $10 \mu\text{s}$ spill in the LE10/185kA (LE250/200kA) beam configuration. About half of these events occur in the calorimeter region and may be fully reconstructed. The Near Detector readout electronics [30] is designed to measure neutrino interactions throughout the spill, without dead-time, and with a timing resolution that allows efficient separation of time-adjacent events. Analog to digital conversion has a floating precision with a minimum least significant bit of 1.4 fC and occurs in contiguous 18.8 ns time-intervals corresponding to the 53 MHz RF of the Main Injector. The electronics has two primary operating modes, “spill-gate” and “dynode”, which are switched in real time. In spill-gate mode the output from every photomultiplier pixel is digitized continuously in a $13 \mu\text{s}$ period starting about $1.5 \mu\text{s}$ before the arrival of neutrinos at the Near Detector. In dynode mode, used for out-of-spill acquisition of cosmic-rays, continuous digitization for a period of 150 ns is initiated independently for each photomultiplier when the dynode signal exceeds a programmable threshold. The integrated trigger rate is typically 30 Hz.

D. Event Reconstruction

The reconstruction procedure uses the topology and timing of hits to identify neutrino interactions inside the detector as well as through-going muon tracks from cosmic-rays or neutrino interactions in the surrounding rock. Two contained vertex interactions are shown in Fig. 10. In ν_μ charged-current events a W boson is exchanged between the neutrino and the target. The final state consists of a muon and hadrons from the recoil system. Neutral-current (NC) events are mediated by Z exchange and only the fraction of energy carried by the recoil system is visible in the final state. The chief goal of the reconstruction procedure is to estimate the visible energy of ν_μ charged-current, ν_e charged-current, and neutral-current interactions while also providing a distilled set of quantities describing the event in order to discriminate between the three processes. As shown in Fig. 10, the strongest evidence of a ν_μ charged-current event comes from the presence of a track that penetrates through several detector planes and is sufficiently distinguishable from any additional hits around the track starting point. These additional hits are associated with the hadronic shower created by the recoil system and their pulse-height may be used to estimate the shower energy. Tracks

¹ This includes interactions occurring throughout the detector, many of which cannot be fully reconstructed but which nevertheless induce some activity in the detector.

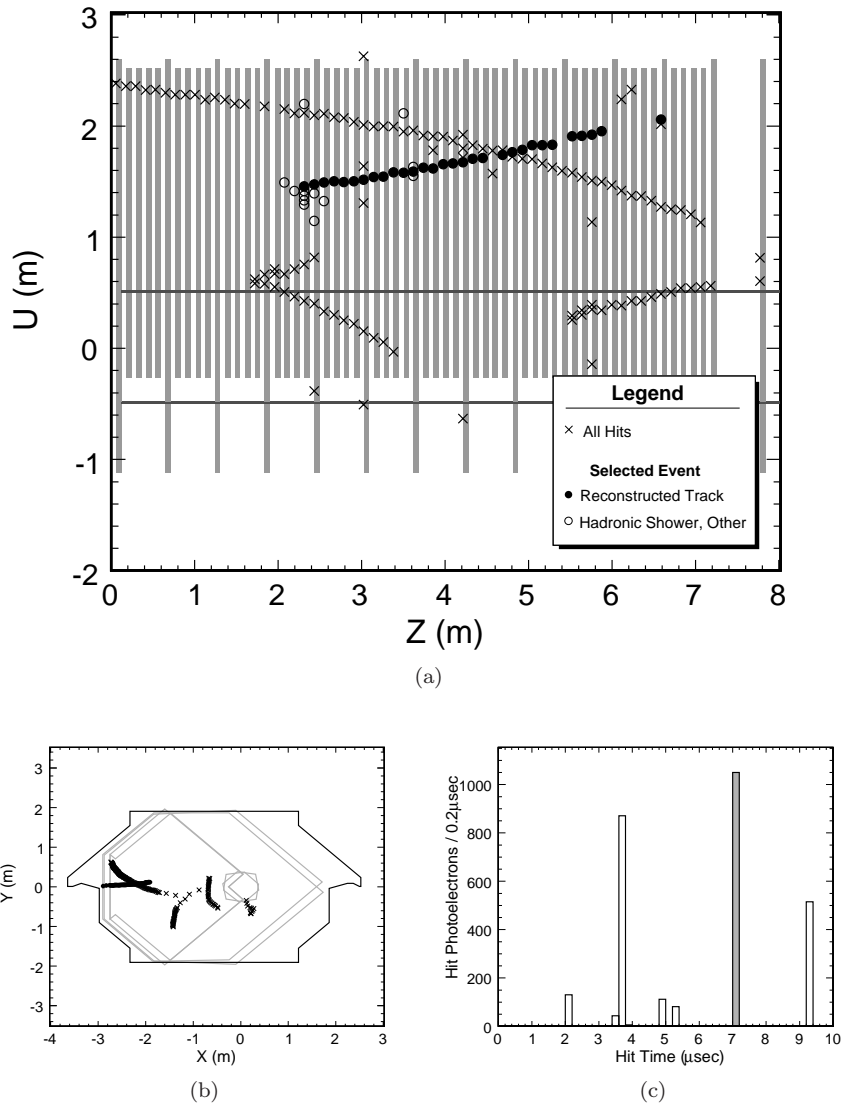


FIG. 11: One beam spill observed in the Near Detector. For clarity of presentation a spill containing a smaller than average number neutrino events was chosen. Data from the one of the two strip orientation is shown in (a). Neutrinos are incident from the left and only the upstream section of the detector is shown. Grey vertical bars indicate the scintillator coverage. The timing and topological pattern of hits in the detector has been used to reconstruct and select an event containing a 5.6 GeV muon and a 2.6 GeV hadronic shower. A beam's eye view of the detector is shown in (b), along with the reconstructed horizontal and vertical positions of track hits in each detector plane. Figure (c) shows the detector signal as a function of time, with signals from the selected event shaded. The bin width is about ten times larger than the detector's timing resolution.

are occasionally found in neutral-current events and additional event-topological quantities are used to discriminate such events from actual ν_μ charged-current events (see Sec. IV A).

The high intensity of the NuMI beam leads to multiple neutrino interactions inside the Near Detector in each beam spill. The first stage in the reconstruction procedure divides the activity in the detector into one or more events, each of which contains hits that are localized in space and time. Figure 11 shows the activity recorded in the Near Detector for a single LE010/185kA beam spill. The plot also indicates how hits from a single interaction can be identified using timing and spatial information. In the Far Detector, where the rate is much lower, there is rarely more than one event per beam spill, and most spills actually contain no neutrino interactions.

A track-finding algorithm is then applied to each event. This algorithm uses a Hough Transform [31] to find track

segments, which are then chained together (taking into account timing and spatial correlations) to form longer tracks. The track momentum is estimated from range² if the track stops within the detector, or from a measurement of its curvature in the toroidal magnetic field if it exits. The curvature measurement is obtained from fitting the trajectory of the track using a Kalman Filter technique, which takes into account bending of the track from both multiple Coulomb scattering and the magnetic field. This procedure also gives an indication of the charge of the reconstructed track. For typical muon tracks produced by beam ν_μ charged-current interactions, the momentum resolution is approximately 5% for range, and 10% for curvature measurements.

Showers are constructed from clusters of strips that are localized in space and time. The energy of a shower is computed from the summed pulse-height of the individual hits, where the pulse-height contribution of any reconstructed tracks that share the same hit is subtracted. The shower energy calibration is discussed in Sec. III F. The energy resolution for neutrino induced hadronic showers is approximately 59% at 1 GeV (32% at 3 GeV).

The total reconstructed energy of each event is estimated by summing the energy of the most energetic track with the energy of any shower present at the upstream end of the track. We select ν_μ charged-current interactions by requiring that the event has at least one well reconstructed track with a starting-point, interpreted as the neutrino interaction point (vertex), in the fiducial volume. In the Near Detector, the fiducial volume is a cylinder of radius 1 m from the beam center³ and length 4 m beginning 1 m downstream of the front face of the detector. The Far Detector fiducial volume is a cylinder of radius 3.7 m from the detector center. Vertices are required to be at least 50 cm from the front and rear planes of the two detector modules and, to assure a track that is long enough to be analyzed, greater than 2 m upstream of the last plane of module two. Tracks are required to have a negative charge to suppress $\bar{\nu}_\mu$ and only events with neutrino energy less than 30 GeV are used so as to preferentially select neutrinos from π , rather than K , decays. A final set of criteria are applied to remove rare periods in which the magnetic field coil was not energized or in which the high voltage was not on in some portion of the detector.

The resulting sample in the Near Detector is 89% pure ν_μ charged-current (91% in the Far Detector, assuming no neutrino oscillations) with the dominant background coming from neutral-current events in which a (usually short) track was reconstructed. Section IV A describes the way in which event-topological quantities were used to improve the sample purity. In the following subsections we describe some characteristics of the dataset before purification.


1. Near Detector Data

We have reconstructed more than 2.6×10^6 neutrino events in the Near Detector in several different beam configurations (see Tab. I). These data have been used to verify and improve the quality of our neutrino interaction simulation, detector calibration, event reconstruction, neutral current rejection procedure (described in Sec. IV A), and beam flux calculation (Sec. V). Since we are ultimately interested in predicting the Far Detector energy spectrum, it is not necessary that the Monte Carlo simulation reproduce the Near Detector data exactly. Instead we must have confidence that any discrepancies have causes which are common to both detectors such that the Near Detector data may be used to improve the Far Detector prediction. For example, uncertainties in the neutrino flux, neutrino cross-sections and energy resolutions (though not the absolute energy scale) may be mitigated by using Near Detector measurements.


The Monte Carlo simulation reproduces many reconstructed quantities in the Near Detector, including the vertex distributions and track angular distribution shown in Fig. 12. Figure 13 shows the mean number of reconstructed events as a function of beam intensity ranging between 5×10^{11} and 2.7×10^{13} POT/spill. The linearity of the curve indicates that the Near Detector is able to measure individual neutrino interactions with little dependence on the interaction rate. The behavior of the curve near the origin suggests a negligible background from non-beam related events. The reconstructed neutrino energy spectrum measured during several months is overlaid in Fig. 14. The average spill intensity in June 2005 is 1.65×10^{13} POT/spill and rises to a maximum of 2.4×10^{13} POT/spill in November 2005. The distributions are consistent within statistical errors and show no rate dependence. We therefore conclude that the Near Detector event reconstruction is stable over this range of spill intensities, and does not introduce any observable biases in the reconstructed neutrino energy spectrum. Consequently, we do not assign any intensity based uncertainty in the detector efficiency.

² When reconstructing the momentum by range, the track is assumed to be due to a muon and the range tables of [32] are employed.


³ The radius is calculated from the beam center in the first detector plane. The cylinders follow the z axis of the detectors which differs from the beam axis by 58 mrad.



(a)



(b)



(c)

FIG. 12: Distributions of basic reconstructed quantities in the Near Detector for events reconstructed and selected using the procedure described in the text. The distributions are normalized to the same numbers of events. The left-hand and center plot show the distribution of reconstructed track vertices for X (perpendicular to detector axis) and Z (parallel to detector axis) detector co-ordinates, and the right-hand plot shows the distribution of the reconstructed track direction relative to the incident neutrino beam.




FIG. 13: Mean number of reconstructed events per Near Detector spill as a function of spill intensity.

2. Far Detector Data

Several criteria were applied to assure the integrity of the Far Detector data. These explicitly exclude data taken during magnet coil and high voltage power failures and periods where the GPS timing system was not operational. The POT-weighted live-time of the Far Detector was 98.9% during the period May 20, 2005 – Feb 25, 2006 in which the oscillation dataset was accumulated. Two additional conditions are applied to Far Detector data in order to reduce contamination from cosmic ray events to a negligible level. First, the direction cosine of the primary reconstructed track with respect to the beam direction must be greater than 0.6 in order to reject the high angle tracks typical of cosmic rays. Furthermore, the event time must lie within a $50 \mu\text{s}$ window around the predicted time of the beam spill at the Far Detector site.

Figure 15 shows the time (relative to the spill time) of the 384 selected neutrino events that satisfy the criteria described above. The width of this distribution is consistent with the Main Injector spill length, and there is no evidence of background contamination within the $50 \mu\text{s}$ window. Figure 16a shows the cumulative distributions of the number of total protons on target and the number of Far Detector neutrino events as a function of time. The two distributions follow each other closely. Figure 16b shows the number of neutrinos per 10^{17} protons on target and exhibits no significant time dependence.

A visual scan of the events was carried out in order to ensure that there was no background contamination in the selected event sample and that the events were well reconstructed. Figure 17 shows the reconstructed vertex distributions of the Far Detector events. The Monte Carlo distributions are normalized to the same number of events as the data to account for the possible effects of oscillations and are in agreement with the data. Additional comparisons between the data and simulation are shown in the sections that follow.




FIG. 14: Reconstructed neutrino energy distributions by calendar month. The distributions are normalized to protons on target and only data obtained in the LE10/185kA configuration is included. Markers representing individual months are offset to clarify the presentation and the line represents the average over all months.



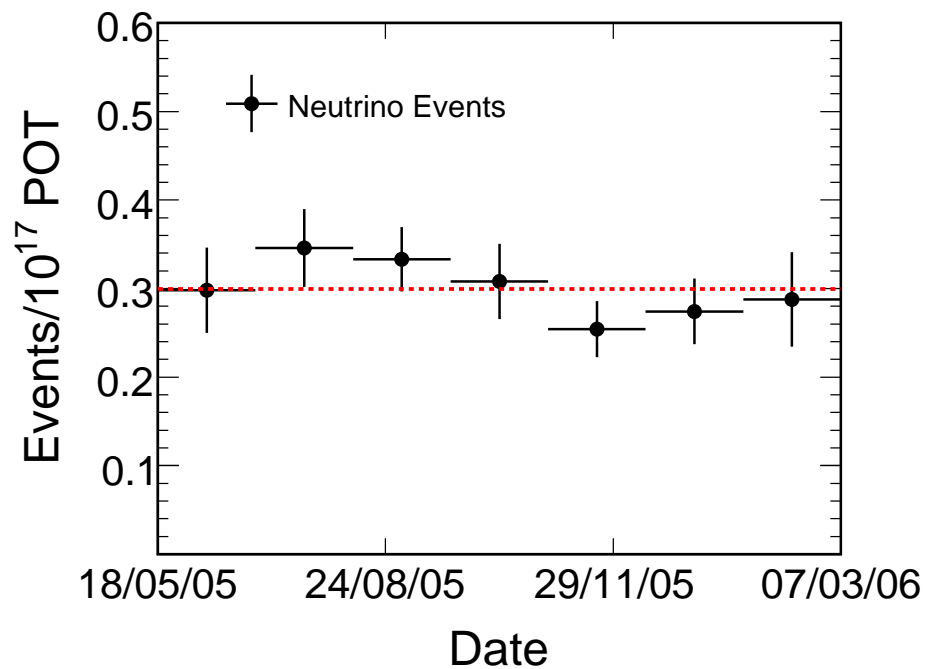


FIG. 15: Distribution of the event time of the 384 Far Detector neutrino events relative to the time of the nearest beam spill [33].




(a)




(b)


FIG. 16: (a) Cumulative distributions of the total number of protons on target and the number of Far Detector neutrino events observed as a function of time. (b) Number of observed Far Detector neutrino interactions per 10^{17} protons on target as a function of time.



(a)



(b)



(c)

FIG. 17: Distributions of track vertices for Far Detector data and Monte Carlo simulation. The Monte Carlo distributions are normalized to the same number of events as the data.




FIG. 18: NEUGEN-v3 calculation of muon-neutrino charged-current cross-sections per nucleon on an isoscalar target. The cross-section per GeV is shown as a function of the neutrino energy for inclusive scattering, quasi-elastic scattering and single pion production. The calculation is compared with experimental data tabulated by [34]. The shaded band corresponds to the cross-section uncertainties described in the text.

E. Modeling Neutrino Interactions

Neutrino interactions are modeled by the NEUGEN-v3 [35] program. NEUGEN simulates both quasi-elastic and inelastic neutrino scattering. The latter includes a Rein-Sehgal [36] based treatment of neutrino induced resonance production, charged- and neutral-current coherent pion production and a modified leading order deep-inelastic scattering (DIS) model [37] extended to improve the treatment in the transition region between DIS and resonant production. KNO scaling [38] is used to calculate the final state multiplicity in the DIS regime. Hadrons produced in the neutrino

scattering are allowed to interact while exiting the target nucleus (“final state interactions”). The final state interaction calculation incorporates pion elastic and inelastic scattering, single charge exchange and absorption [39]. The calculation is benchmarked by a comparison of final states in $\nu + d$ and $\nu + \text{Ne}$ interactions as measured in the BEBC and ANL–12 ft bubble chambers [40].

Figure 18 shows the ν_μ charged-current cross-section as a function of neutrino energy in the laboratory frame. Based on a comparison of the model predictions to independent data, some of which [34] is shown in Fig. 18, we assign a systematic uncertainty of 3% on the normalization of the DIS ($W > 1.7 \text{ GeV}/c^2$) cross-section, and a 10% uncertainty in the normalization of the single-pion and quasi-elastic cross-sections. We estimate a 20% uncertainty in the relative contribution of non-resonant states to the 1π and 2π production cross-sections for $W < 1.7 \text{ GeV}/c^2$. This uncertainty was determined from the parameter uncertainties and variations observed in fits to both inclusive and exclusive channel data, and in fits to data in different invariant mass regions. Final state interactions are expected to have a significant effect on the visible energy of the hadronic final state [41]. In particular there are significant uncertainties in the rate of pion absorption, the mechanism for transferring the pion’s energy to a nucleon cluster, and the amount of energy eventually visible to the detector. We account for these uncertainties by studying the shift in the reconstructed shower energy when we turn the simulation of final-state interactions off, and when we modify the simulation so that all of an absorbed pion’s energy is lost. We find that the predicted response to hadronic showers changes by approximately 10% [41] in these two extreme cases and use this as a conservative estimate of the uncertainty on the absolute hadronic energy scale.

The MINOS detector simulation is based on GEANT3 [25] and is used to generate raw energy depositions (GEANT hits) which serve as the input to our detector response model. The simulation randomly samples neutrinos from the flux predicted by the beam simulation (Sec. II B) and traces them through the Near and Far Detector halls. Events are generated inside the detectors as well as in the surrounding support structure and rock. The simulation includes a detailed geometric model of the detector which describes the material crossed by neutrinos and neutrino induced tracks to within 1% plane-to-plane and 0.3% averaged over the detector. The position of individual scintillator strips was determined with a precision of approximately 1 mm using cosmic-ray tracks. The magnetic field is modeled via finite element analysis driven by bench measurements of the steel B–H curve. Since performing this analysis we have recalibrated our field and found that it increased by 12.3% and 9.2% averaged over the fiducial volumes in the Near and Far Detectors, respectively. This recalibrated field shifts the momentum scale for muons exiting the Near Detector by approximately 6.2% (4.6% in the Far Detector) but does not significantly affect the scale for muons which stop in the detector.

The detector response simulation has been tuned by directly incorporating the channel-by-channel calibration constants determined by the detector calibration procedure (see Sec. III F below). Additional data, measured in pre-installation tests, were used to constrain details of the scintillator and photomultiplier response models. The simulation has been benchmarked against through-going and stopping cosmic rays in the Near and Far Detectors, beam neutrino events in the Near Detector, as well as a series of test-beam measurements [42, 43, 44, 45] collected with a scaled down version of the MINOS detectors in the CERN PS East Hall. The test-beam measurements are used to fix the energy scale and validate the simulation of electromagnetic and hadronic showers. The hadronic shower code GCALOR [46] was found to be in relatively good agreement with the data (see Sec. III F) and is used in our simulations.

F. Detector Calibration

The principal tools for calibrating the detector response are an LED based light-injection (LI) system, a test-bench scan of the scintillator modules with a radioactive source and cosmic-ray muons. The detector is calibrated in a multi-stage procedure that converts the raw photomultiplier signal $Q_{raw}(i, t, x)$ measured by channel i at time t for an energy deposition at position x into a fully corrected signal Q_{cor} . Each calibration stage produces a “calibration constant”. The fully corrected Q_{cor} is defined as the product of $Q_{raw}(i, t, x)$ and the calibration constant from each stage:

$$Q_{cor} = Q_{raw}(i, t, x) \times D(t) \times U(i) \times A(i, x) \times E$$

where D, U, A and E refer to:

Drift correction $D(t)$: The channel gains and their variation over time are measured with the LED based light-injection system, demonstrating that short-term ($< 24 \text{ h}$) gain variations are small and occurred mostly due to environmental changes in the detector halls. Light-injection data were eventually superseded by measurements of the mean signal per plane induced by through-going cosmic-ray muons since the muon data were also able to correct for variations in the scintillator light-output. The detector response varies by $< 2\%$ over the data-taking period as shown in Fig. 19. The decreasing response in the Far Detector is likely due to aging of the scintillator.




FIG. 19: Variations in the median signal per plane deposited by through-going cosmic-ray muons observed during the data-taking period covered by this paper. The time dependence is largely due to variations in the environmental conditions in the Near and Far Detector halls and aging of the scintillator. The zero point on the ordinate is arbitrary.

Uniformity correction $U(i)$: Through-going cosmic-ray muons are used to account for differences in light output between individual strips as well as attenuation in the optical fibers and connections. Event-by-event corrections are applied to account for the muon track angle and the expected inefficiency due to statistical fluctuations at low light-levels. The calibration reduces the uncorrected 30% strip-to-strip response variations to approximately 8%.

Attenuation correction $A(i, x)$: A radioactive source was used to map out the response of each scintillator module at many positions along each strip. This was done on a test-bench setup prior to installation of the modules in the Near and Far Detectors. The data were then fit to an empirical model of optical attenuation in the wavelength-shifting and clear optical fibers. The resulting parametrizations are used to correct the signals from cosmic-ray muons during the calculation of the uniformity calibration constants $U(i)$ and also to correct the reconstructed shower energy for attenuation based on the reconstructed shower position. The relative size of the correction varies by about 30% over the 8 m length of a Far Detector scintillator strip when signals from both ends are added, and by about 50% over the 3 m length of a Near Detector strip.

Signal scale calibration E : The overall scale of the signals is established by the detector's response to stopping muons. This provides the standard which fixes the absolute calibration of the Near and Far Detectors, allowing




FIG. 20: Mean signal per scintillator strip induced by beam muons. Corrections are made with cosmic-ray muons and account for non-uniformities in channel gain $D(t)$, attenuation in optical fibers $A(i, x)$, and strip light output $U(i)$.

their signals to be compared to one another. The calibration is done by tabulating the detector's response to muon crossings using only the portion of each track in which the muon energy is between 0.5 GeV and 1.1 GeV, deduced from the distance to the track endpoint. This energy window avoids the rapid variation in dE/dx near the track's end. Prior to this calibration the signals have already been corrected for D, A, U as described above. For each muon a correction is applied to account for the muon's path length in each scintillator plane. The mean response is then calculated for each individual strip and a single constant representing the entire detector is derived from the median over all strips.

The effect of the uniformity and attenuation calibration stages is shown in Fig. 20. After calibration, the fully corrected signal is expressed in muon equivalent units (MEU) as defined by the signal scale (E) calibration procedure. One MEU corresponds to approximately 3.7 photoelectrons at the center of the Far Detector (7.5 photoelectrons when signals from both strip-ends are summed) and 5.4 photoelectrons at the beam center in the Near Detector. Based on muon stopping-power tables [32] we find that one MEU corresponds to 2.00 ± 0.02 MeV of muon energy loss in scintillator [44, 47]. This calibration was independently derived at the Near and Far Detectors with a relative

uncertainty of $\pm 2\%$.

Data from a dedicated, test-beam calibration detector [47] are used to benchmark the hadronic and electromagnetic shower simulation done by `GEANT/GCALOR` [42, 43]. The measurements demonstrate that the simulated response agrees with the data at the level of 1–5%, depending on the energy and particle type. This level of agreement is inclusive of energy dependent offsets in the beam momentum of the test beam-lines and uncertainties in the simulation of energy loss upstream of the calibration detector. In light of the much larger uncertainty associated with final state interactions, no attempt was made to resolve these discrepancies when determining the conversion from MEU to GeV for neutrino induced showers. Instead a conservative uncertainty of 5.6% is assigned to the absolute shower energy scale to account for the data/MC discrepancies from the calibration detector as well as the precision of the stopping muon calibration performed there. An energy dependent MEU to GeV conversion is then extracted from the simulation such that the reconstructed shower energy estimates the energy transferred to the hadronic system.

IV. CC EVENT SELECTION

A. Event Classification

In the MINOS experiment, neutrino oscillations are expected to cause a deficit of ν_μ charged-current events at energies $E_\nu \lesssim 5$ GeV for $\Delta m^2 \approx 2.5 \times 10^{-3} \text{ eV}^2/c^4$ and our baseline of 735 km. Neutral-current events which have been misidentified as ν_μ charged-current tend to populate this energy range and could obscure the oscillation signal. We remove neutral-current events from the oscillation sample with a technique based on the event topology. The technique uses three probability density functions (PDFs), which are constructed for the following variables: (a) the event length, expressed in units of the number of detector planes, (b) the fraction of the total event signal in the reconstructed track, and (c) the average signal per plane induced by the reconstructed track. These quantities are related to the muon-range, the event inelasticity, and the average energy loss dE/dx of the muon track and are distributed differently for ν_μ charged-current and neutral-current events as shown in Fig. 21. The probability that a particular event is consistent with the ν_μ charged-current or neutral-current PDFs is given by the product of the three individual probabilities

$$P_{CC,NC} = \prod_{i=1}^3 f_i(x_i)_{CC,NC} \quad , \quad (4)$$

where the $f_i(x_i)$ are the individual PDFs for ν_μ charged-current and neutral-current events respectively. The PDF distributions are well modeled by the simulation in the regions in which the neutral-current and ν_μ charged-current samples overlap and for neutrino energies above approximately 10 GeV where the Monte Carlo simulation is better constrained by data from previous experiments and the neutrino flux has a weaker dependence on the energy.

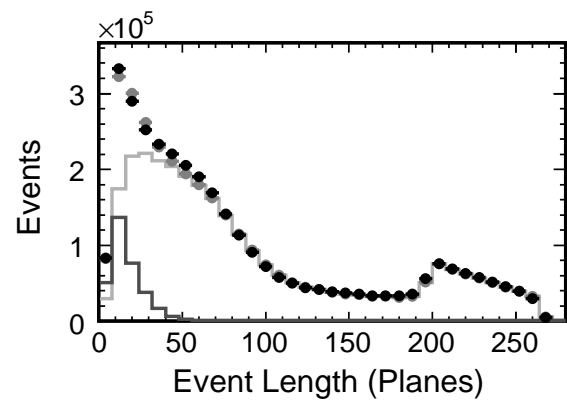
An event selection parameter S is derived from P_{CC} and P_{NC} according to

$$S = - \left(\sqrt{-\ln(P_{CC})} - \sqrt{-\ln(P_{NC})} \right) \quad . \quad (5)$$


Events that are more likely to originate from the ν_μ charged-current interactions give positive values, and those that are more likely to be neutral-current give negative values. Figure 22 shows the distribution of S for data collected in the Near Detector in the LE10/185kA, LE100/200kA and LE250/200kA beam configurations. The behavior of the data is well-reproduced by the Monte Carlo simulation especially in the region $S > -0.1$ in which ν_μ charged-current events dominate.

We isolate an enriched sample of ν_μ charged-current events for use in the oscillation analysis by requiring $S > -0.2$ in the Far Detector and $S > -0.1$ in the Near Detector. These values were chosen to optimize the statistical sensitivity to the oscillation parameters Δm^2 and $\sin^2 2\theta$ in the Far Detector, and to provide a selected event sample in the Near Detector that has the same purity, defined as the fraction of selected events that are true ν_μ charged-current interactions, as the Far Detector. Figure 23 shows the selection efficiency, defined as the fraction of true ν_μ charged-current events that have $S > -0.1$ in the Near Detector ($S > -0.2$ in the Far Detector), and the resulting purity of the sample in the two detectors. The efficiency is calculated with respect to the events passing the criteria of Sec. III D and varies slowly as a function of reconstructed energy above 0.5 GeV. The contamination from misidentified neutral-current interactions is greater at low reconstructed energy, hence the purity of the selected sample drops below 1 GeV. The Monte Carlo simulation indicates that the average selection efficiency and purity (integrated over the unoscillated Far Detector LE010/185kA spectrum) are 90.0% and 98.2% respectively.

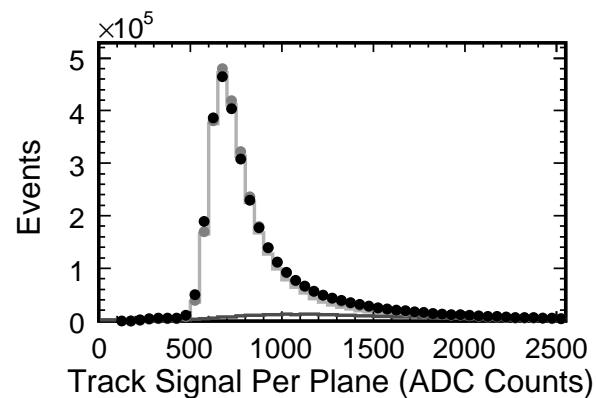
The PDFs used to discriminate between ν_μ charged-current and neutral-current interactions depend on the ν_μ energy spectrum and are therefore sensitive to neutrino oscillations. As such, the discrimination technique was



(a)



(b)



(c)

FIG. 21: Distributions of the variables used to define the event separation parameter S in Eq. 4 and Eq. 5 for events satisfying the criteria of Sec. III D. Near Detector data collected in the LE10/185kA beam configuration is shown.

Selection Criterion	Far Detector Events
Reco. track in fiducial vol.	427
Data quality cuts	408
Event Timing cut	404
Beam Quality cuts	390
Track direction cut	384
Track quality cuts	365
Negative track charge	306
Reco. energy < 30 GeV	275
$S > -0.2$	215

TABLE II: Effect of selection criteria on the Far Detector neutrino event sample. The first eight criteria are described in Sec. III D. In accordance with our analysis strategy, the total number of events passing all cuts was not revealed until all of the data quality checks were complete and our analysis method was fully defined.

developed and validated by comparing the Monte Carlo simulation to data collected in the Near Detector. A fraction of the Far Detector dataset was available for data quality checks. Examination of oscillation-sensitive distributions in the full Far Detector data set was only performed after these checks had been made and we had fully defined our analysis procedure. Table II shows the effect of our selection on those Far Detector events that are recorded in coincidence with NuMI beam spills. After requiring $S > -0.2$ we retain 215 of the 275 events passing the criteria of Sec. III D. The variables used in Eq. 4 and Eq. 5 are shown in Fig. 24 for the 275 event sample. Two versions of the simulation are shown. The first does not account for neutrino oscillations while in the second we apply the effect of oscillations according to the mixing parameters extracted from our data using the procedure described in Sec. VII. The measured deficit of events with respect to the unoscillated simulation is consistent with the partial disappearance of ν_μ charged-current events at energies less than about 10 GeV and is well described when we account for oscillations.

B. Backgrounds

1. Neutral-Current Contamination

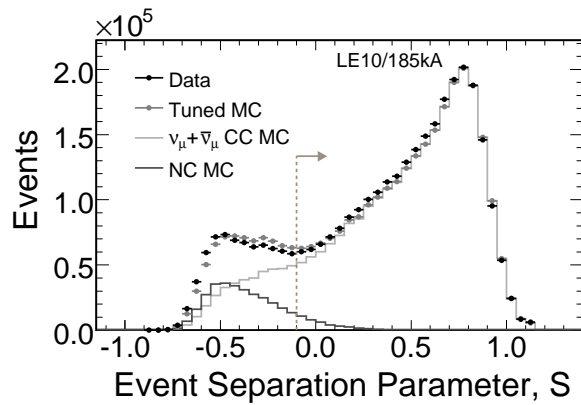
The neutral-current background in the Far Detector is estimated from the Monte Carlo simulation as the number of neutral-current events that pass the event selection cut. These events comprise only $< 2\%$ of the selected unoscillated ν_μ CC sample but tend to congregate in the lowest reconstructed energy bins, as shown in Fig. 23, and so constitute a significant background to the ν_μ charged-current oscillation signal. In this section, we describe how we have used Near Detector data to directly provide an estimate of the systematic error on the number of neutral-current events that are selected in the final Far Detector event sample.

To handle uncertainties that arise from the modeling of hadronic showers in the Monte Carlo, showers from cleanly identified ν_μ charged-current events in the Near Detector with the muon tracks removed were used to model neutral current events. The ratio between the track-removed shower S distributions for data/MC provides a S -dependent Monte Carlo scaling factor for the shape of the neutral-current event selection parameter distribution.

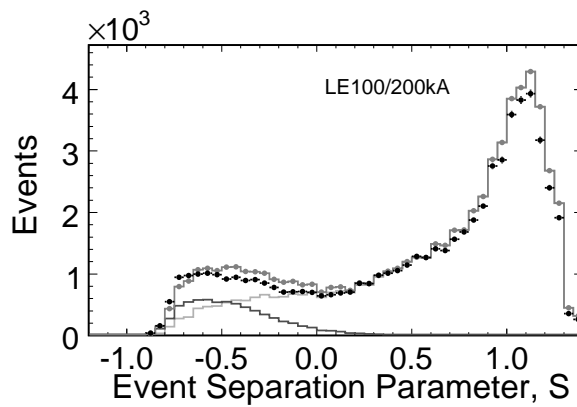
To determine the uncertainty on the overall neutral-current signal normalization, the event selection parameter distribution S in the data was fitted to determine the amplitude of the ν_μ charged-current and neutral-current signals. In this fit, the shape of the charged-current component was obtained from the Monte Carlo simulation, while the neutral-current shape was scaled according to the data/MC ratio observed in the track-removed ν_μ charged-current event sample. The fits to the data for the signal amplitude and the data/MC comparison of the track-removed ν_μ charged-current events were performed in six reconstructed energy slices.

A comparison of the event separation distribution S for the simulated neutral-current and selected muon removed ν_μ charged-current events yield reasonable agreement, as shown in Fig. 25. However, the muon removed ν_μ charged-current events from Near Detector data, which are represented by the black points in Fig. 25, peak at lower values of S than the corresponding Monte Carlo distribution. This discrepancy implies that the data contain shorter events and have higher average pulse height per plane than the Monte Carlo events.


The ratio of the data to simulated S distributions obtained from the muon-removed events is used to provide a reweighting of the shape of the S distribution for neutral-current Monte Carlo events. This procedure is performed in six bins of reconstructed energy and in all cases, the reweighted distributions predict fewer neutral-current events



(a)




(b)




(c)

FIG. 22: The event separation parameter S , as defined in Eq. 5, plotted for Near Detector data and simulation, and for three beam configurations in Tab. I.



(a)



(b)

FIG. 23: Efficiency for identifying ν_μ charged-current events using the event separation parameter defined in Eq. 5, along with the purity of the resulting sample. The efficiency is defined with respect to the events satisfying the criteria of Sec. III D.




FIG. 24: Distributions of the variables that are used in ν_μ charged-current and neutral-current event classification, for Far Detector data and Monte Carlo simulation. The bottom-right plot shows the distribution of the event separation parameter S that is derived from these three PDFs. Oscillations are accounted for according to the mixing parameters extracted from our data (see Sec. VII). These distributions are shown here to illustrate the level of agreement between data and simulation for S and its input variables. They were not examined until the analysis procedure had been fully defined.

for $S > -0.1$ than the *a priori* simulation, by an amount that ranges from 20-50%, as shown in Table III.

The origin of this difference is most likely due to an unknown combination of shower modeling and/or neutral-current cross-section uncertainties. We have thus assigned 50%, the largest difference observed in the fit above $S = -0.1$, as the error on the neutral-current background.

2. Cosmic Ray Background

We have estimated the background rate from cosmic-ray muons in the Far Detector selected sample using two independent methods. First, we examined the rate of selected events in the spill trigger window that are not within the expected $10 \mu\text{s}$ wide window around the time of true beam spill (see Fig 15). Secondly, we applied the ν_μ charged-current selection criteria described above to a sample of 2.3×10^6 spill triggers taken in anti-coincidence with the beam spill. Both of these methods yielded an upper limit (90% C.L) of 0.5 background events in a $10 \mu\text{s}$ wide window around the time of true beam spill for an exposure of 1.27×10^{20} POT. The background due to neutrino-induced ν_μ

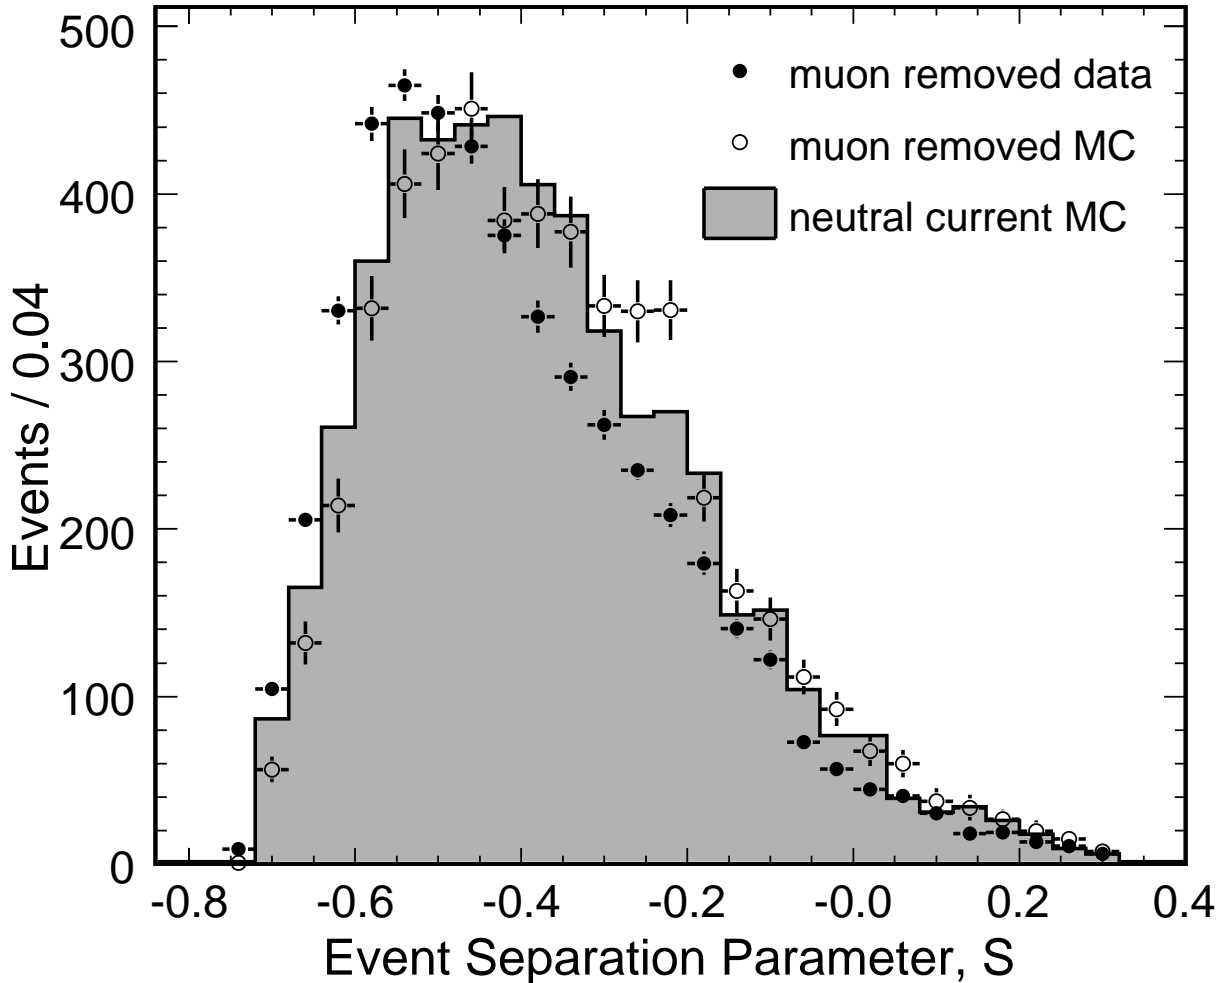


FIG. 25: Distribution of the event separation parameter S for true neutral-current Monte Carlo events (shaded histogram), for selected ν_μ Monte Carlo events with the reconstructed muon track removed (open circles) and for selected Near Detector data events with the muon track removed (filled circles).

charged-current interactions in the rock upstream of the Far Detector is estimated from the Monte Carlo simulation to be 0.38 events. The rate of cosmic-ray muons is about 30 times larger in the Near Detector than it is in the Far Detector. The cosmic-ray background is negligible, however, due to the much greater ($\times 10^6$) neutrino flux at the Near Detector.

V. CONSTRAINTS ON THE NEUTRINO FLUX FROM NEAR DETECTOR DATA

Figure 26 compares the measured ν_μ charged-current energy spectrum with the Monte Carlo prediction for six different beam configurations. There are noticeable differences between data and the Monte Carlo calculation in all configurations but the magnitude of the discrepancies and the energy range over which they occur depends on the beam configuration. These observations suggest that a significant source of the disagreement between data and the Monte Carlo simulation may be due to inaccuracies in the calculation of the neutrino flux rather than mis-modeling of neutrino interactions or detector acceptance, since the latter depend most strongly on the energy of the incident neutrino while the former depends on the beam configuration.

As noted in Sec IID uncertainties in the neutrino flux calculation arise from insufficient knowledge of hadron production off the NuMI target as well as from several beam focusing effects described in Sec. IID. We constrain our

Energy Range	Data/MC ratio
0-1 GeV	0.642
1-2 GeV	0.503
2-4 GeV	0.804
4-6 GeV	0.583
6-12 GeV	0.434
12-30 GeV	0.579

TABLE III: The ratio of the number of neutral-current events with $S > -0.1$ in the simulation and the best fit to the data using the neutral-current shape derived from the track-removed ν_μ charged-current events. The data indicates that the neutral-current background is over-predicted by the simulation.

parameter	initial uncertainty	best fit
Horn 1 misalignment	1 mm	0.1 ± 0.4 mm
POT Normalization	2%	$0.0 \pm 0.8\%$
Horn Current Distribution k	0.17 mm^{-1}	$0.07 \pm 0.05 \text{ mm}^{-1}$
Shower Energy Offset	100 MeV	24.7 ± 100.0 MeV
Baffle scraping	0.25%	$0.1 \pm 0.1\%$
Horn current calibration	1%	$0.9 \pm 0.3\%$
Neutrino Energy Mis-calibration	10%	$-4.9 \pm 10.0\%$
Neutral-Current Contamination	30%	$-14 \pm 7\%$

TABLE IV: Detector and beam modeling effects included when fitting the Monte Carlo to the Near Detector data. The effect of a 1σ variation in each of the beam focusing parameters is shown in Fig. 7a. The horn current distribution as a function of depth z in the conductor is modeled as $I(z) \propto \exp(-kz)$ and we tabulate k here.

flux calculation using ν_μ charged-current energy spectra measured in the Near Detector. The uncertainty in hadron production can be constrained because the position of the target and the horn current determine the region of pion (p_z, p_T) which contributes to ν_μ at the Near Detector. This is shown in Fig. 27 for each of the beam configurations in which we collected data (see Tab. I). Our method works by representing the underlying production yield, $d^2N/dp_z dp_T$, as a parametric function which we use to tune the Monte Carlo in a χ^2 fit to the Near Detector data. We add terms to the χ^2 which describe the influence that beam focusing and detector modeling uncertainties have on the ν_μ spectrum (see Fig. 7a). This technique is similar to those used in previous experiments [51, 52, 53, 54, 55] but the multiple beam configurations allow us to selectively enhance different regions of p_z and p_T .

The FLUKA05 prediction of π^+ yields off the NuMI target is well described by the function

$$f(p_z, p_T) = \frac{d^2N}{dp_z dp_T} = [A(p_z) + B(p_z)p_T] \exp\left(-C(p_z)p_T^{3/2}\right) \quad (6)$$

where A , B and C are functions of p_z . Equation 6 and A , B and C are similar to the functions advocated by [50] but we have modified them to better describe the thick target yields. The ν_μ spectra are then fit by warping A , B , C as linear functions of p_z . The fit outputs the warped functions A' , B' , C' and we calculate hadron yield scale factors relative to FLUKA05 as

$$w(p_z, p_T) = \frac{f(p_z, p_T; A', B', C')}{f(p_z, p_T; A, B, C)} \quad (7)$$

These scale factors are shown in Fig. 28.

In our fit we include a term to constrain the mean transverse momentum of pions, $\langle p_T \rangle$, to the FLUKA05 value of $364 \text{ MeV}/c$, with an uncertainty of $15 \text{ MeV}/c$, obtained from the variation of independent hadron production calculations [24, 48, 49, 56, 57]. The K^+ yield does not contribute significantly to the ν_μ flux below 30 GeV and was permitted to vary by a scale factor. We augment the fit to include terms which represent distortions of the ν_μ charged-current spectra caused by uncertainties in the horn current, horn position, distribution of current in the horn conductors, and proton beam scraping on the baffle. The effect of these uncertainties are shown in Fig. 7a. We also include a 2% uncertainty in the normalization due to proton on target counting. To avoid forcing the fit to over-correct for detector acceptance uncertainties we include a 10% error on the neutrino energy, a 100 MeV uncertainty in the shower energy and a 30% uncertainty in the normalization of the neutral-current contamination.

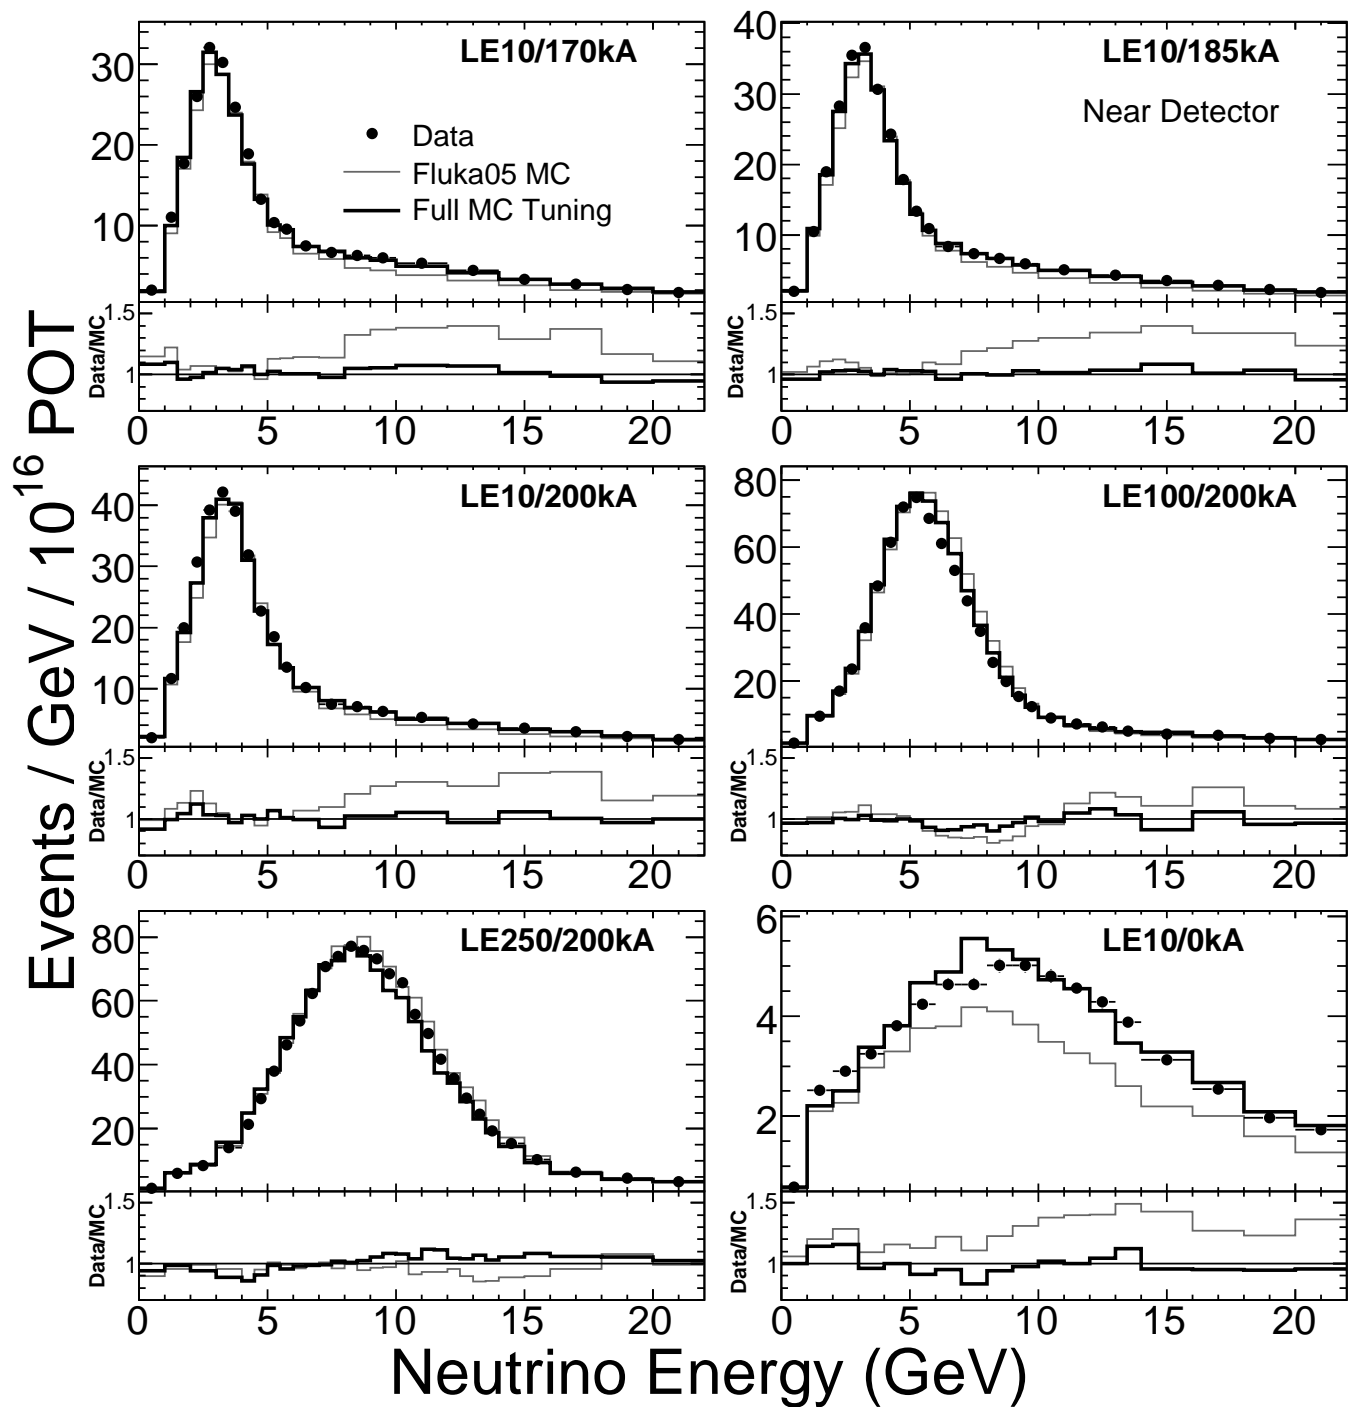


FIG. 26: ν_μ charged-current energy spectra measured in the six beam configurations of Tab. I and compared with the Monte Carlo prediction. Two Monte Carlo predictions are shown: one (thin line) with the *ab initio* calculation based on FLUKA05, the other (thick line) after constraining hadron production, focusing and detector parameters with the neutrino data. Panels along the bottom of each figure show the ratio of the measured and simulated spectra.




FIG. 27: The yield, $d^2N/dp_z dp_T$, of π^+ off the target that contribute to the ν_μ charged-current event rate in each of the beam configurations of Tab. I. Here p_T and p_z refer to the transverse and longitudinal momentum of pions leaving the NuMI target before entering the focusing horns. The box areas are proportional to the probability that a pion has the given (p_z, p_T) and results in a ν_μ charged-current interaction in the Far Detector. As is evident, each beam configuration samples a different region of (p_z, p_T) .




FIG. 28: Pion production scale factors, as a function of p_z, p_T , derived by fitting Eq. 6 to the Near Detector data according to the procedure described in the text. The scale factors are relative to the FLUKA05 prediction.

The results of the fit are shown in Fig. 26, Fig. 28 and Tab. IV. The agreement between the measured and calculated neutrino energy spectra in Fig. 26 is significantly improved with residual discrepancies reduced to 5% or less in all beam configurations. The χ^2 per degree of freedom improves from 3087/360 to 569/345. Figure 28 shows the scale factors applied to the hadron production yield $d^2N/dp_z dp_T$. The most significant adjustment is made to the region $p_T \lesssim 150 \text{ MeV}/c$ and $p_z \gtrsim 15 \text{ GeV}/c$ to increase the flux of unfocused hadrons which dominate the flux for neutrino energies larger than 8 GeV in the LE10 beam configurations. Only modest adjustments are made in the region focused in the LE10 configurations ($100 \lesssim p_T \lesssim 500 \text{ MeV}/c$, $p_z \lesssim 10 \text{ GeV}/c$). Hadrons produced with high p_T and p_z do not contribute to neutrinos in the MINOS detectors and the scale factors are not constrained. The best fit parameters, shown in Tab. IV, are all within 1σ of their initial values. Figure 29 shows the predicted ratio Far/Near of ν_μ charged-current energy spectra before and after the fitting procedure. The fit improves our understanding of the relationship between the Far and Near spectra and we use the results in all further Monte Carlo calculations. Uncertainties in the ratio are shown as boxes which are reduced after the fit as a result of the constraints the neutrino data places on the functions A, B, C of Eq. 6 and the parameters of Tab. IV. We have checked that alternate parametrizations of $d^2N/dp_z dp_T$ yield similar Far/Near ratios. This is important since the ratio characterizes the way in which the neutrino flux changes between the two detectors and indeed provides a way in which the Near Detector data can be used to predict the Far.




FIG. 29: The predicted ratio (Far/Near) of ν_μ flux at the two detectors before and after the tuning procedure described in the text. The height of the box around each point denotes the uncertainty on the ratio. For the points labeled “Fluka 2005” the uncertainties were calculated from the *a priori* uncertainties shown in Fig. 7b as well as the range of predictions from several hadron production codes [24, 48, 49, 50]. For the points labeled “Tuned MC” the uncertainties are computed from the post-fit uncertainties on the functions A' , B' and C' of Eq. 6 as well as those on the additional parameters in Tab. IV. Unless specified otherwise, the tuned Monte Carlo is used in all further calculations.

VI. PREDICTION OF THE FAR DETECTOR SPECTRUM

The Near Detector measures the neutrino energy spectrum close to the source and before oscillations have occurred so as to mitigate the significant uncertainties in the neutrino flux, cross-sections and detector acceptance. As discussed in Sec. V these data have been used to improve the flux calculation. The primary objective, however, is not to constrain the flux, cross-section or acceptance independently but instead to predict the energy spectrum at the Far Detector in the absence of neutrino oscillations. Any differences between the prediction and the Far Detector data may then be interpreted as neutrino oscillations or some other hypothesis.

A. Extrapolation Techniques

We refer to the process of predicting the Far Detector spectrum as “extrapolation” because of the crucial role played by the Near Detector data, and the expected similarity⁴ between the neutrino fluxes at the two sites (see Fig. 5, Fig. 29 and the discussion in Sec. II C). The extrapolation process may then, in one approach, be viewed as making $\sim 30\%$ ⁵ adjustments based on the energy spectrum measured in the Near Detector in order to predict the Far Detector spectrum. These adjustments are made, in part, by using the Monte Carlo simulation but are relative between the two detectors and are thus less sensitive to uncertainties in the absolute flux, cross-sections and detector acceptance. We refer to this approach as “direct” extrapolation because the Near Detector data are used as measured without further constraining the simulation. In a second approach Near to Far Detector extrapolation may be viewed as using the Near Detector data to constrain the Monte Carlo calculation of the neutrino flux, cross-sections and detector acceptance. The improved simulation is then used to calculate the energy spectrum expected at the Far Detector. This approach is referred to as “indirect” extrapolation in contrast to the “direct” approach. In practice the two approaches are complementary and result in very similar predictions of the Far Detector spectrum. The present section describes the specific techniques that were developed to implement each approach. We have pursued a total of four techniques in order to better understand the extrapolation process and study the robustness of our oscillation measurement.

1. Direct Extrapolation Methods

In the absence of oscillations the neutrino fluxes at the Far and Near Detectors are similar, but not identical, and the differences can be most transparently characterized in terms of the Far to Near flux ratio shown in Fig. 29. The energy dependence in the ratio comes from two main sources [58, 59]. First, the solid angle acceptance of the Near Detector varies as a function of the distance from the neutrino production point while the acceptance of the Far Detector is essentially constant. The two detectors subtend different solid angles and by Eq. 3 observe different neutrino angle distributions which, according to Eq. 2, result in different neutrino energy distributions. Apertures in the beam-line (for example, the decay pipe) and under or over focusing in the horns enhance this effect by attenuating pions in certain energy ranges and locations along the beam-line. According to Eq. 2, this introduces an energy dependent suppression of the neutrino flux which, because of the disparate angular acceptance, is not the same in the two detectors. Second, neutrinos produced by decays at finite radii from the beam axis will cover different angular ranges when intersecting the fiducial volumes of the Near and Far Detectors. This effect also introduces different energy distributions in the two detectors. These effects cause the peaks and dips in Fig 29. Thus, the major sources of *relative* Near to Far differences in the neutrino flux (and hence, energy spectrum) are largely due to beam-line geometry, focusing, angular acceptance and decay kinematics. This suggests that the Monte Carlo simulation may be used, with relatively small uncertainties, to derive a transfer function that extrapolates the neutrino energy spectrum measured in the Near Detector to the Far Detector.

The Far to Near flux ratio of Fig. 29 is itself nearly a transfer function except that it is expressed in true neutrino energy and does not account for detector energy resolution, small disparities in efficiency between the two detectors, and their different fiducial masses. A suitable replacement, however, can be evaluated by applying the ν_μ charged-current selection criteria of Sec. IV A to fully simulated events in both detectors to derive neutrino event rates, $\{n_i\}$ and $\{f_i\}$ in bins i of reconstructed neutrino energy. The Near Detector data $\{N_i\}$ are then used to predict the Far spectrum:

$$F_i^{\text{predicted}} = N_i \times \frac{f_i}{n_i} \quad \text{“The F/N Method”} \quad . \quad (8)$$

This technique is referred to as the “F/N method” but is equivalent to scaling each bin in the simulated Far Detector reconstructed neutrino energy spectrum by the ratio of the number of observed to expected events in the corresponding Near Detector reconstructed neutrino energy bin. The effect of neutrino oscillations may be accounted for by modifying f_i according to Eq. 1.

Neutrinos having a given energy in the Near Detector come from decays which would, collectively, yield neutrinos covering a range of energies in the Far Detector as indicated in Fig. 5, Eq. 2 and Eq. 3. This effect is shown at several

⁴ Modulo the ratio of solid angles $\Omega_{FD}/\Omega_{ND} = d_{ND}^2/d_{FD}^2 \approx (1 \text{ km}/735 \text{ km})^2$ where d_{ND}, d_{FD} refer to the distance from the target to the detectors.

⁵ These adjustments are relative to the solid angle correction Ω_{FD}/Ω_{ND} .

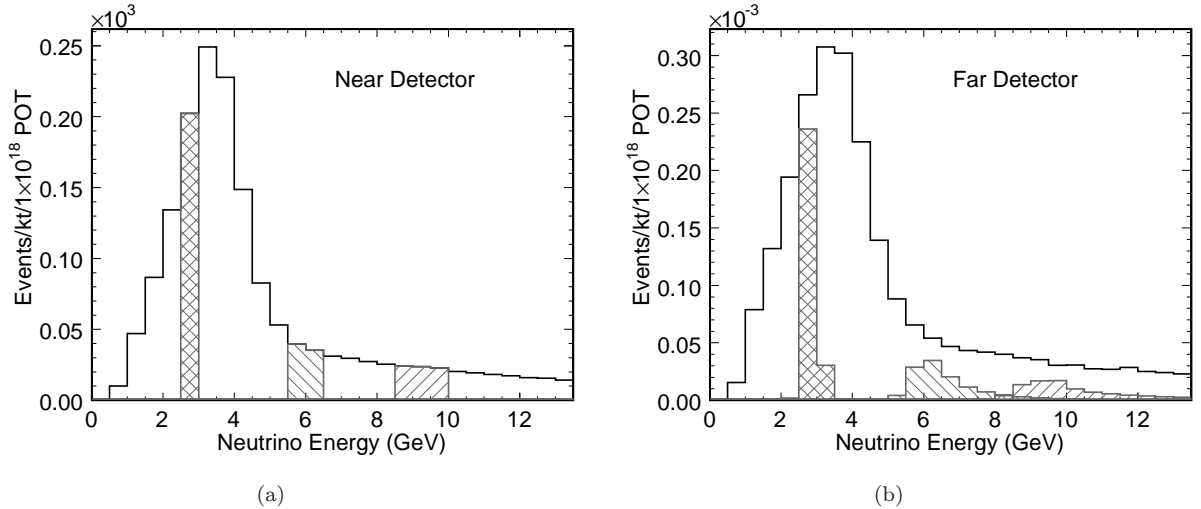


FIG. 30: The relationship between the energy of neutrinos observed in the Near Detector with those observed in the Far Detector. Decays producing neutrinos with a given energy in the Near Detector would produce a range of energies in the Far Detector, yielding the energy smearing seen here. This relationship may be represented algebraically by treating the joint distribution in Fig. 31 as a matrix, and the distributions (a) and (b) as column vectors as in Eq. 9.

energies in Fig. 30 and is again a consequence of the differences in angular acceptance of the two detectors, beam-line apertures, decay kinematics and off-axis decays. The fact that a single energy in the Near Detector corresponds to a range in the Far Detector suggests that the neutrino energy spectra may be related by a two-dimensional matrix rather than a one dimensional ratio. This “beam matrix” $\{B_{ij}\}$ is shown in Fig. 31. Each cell represents the number of neutrinos expected in energy bin i at the Far Detector for one neutrino in bin j in the Near Detector. Neutrinos with energies in the range 1-6 GeV are well focused and the contents of the cell along the diagonal is approximately proportional to the ratio of solid angles $\Omega_{FD}/\Omega_{ND} = d_{ND}^2/d_{FD}^2 \approx (1 \text{ km}/735 \text{ km})^2 \approx 1.8 \times 10^{-6}$. The matrix is constructed from the beam simulation using the known geometric acceptance of the two detectors. The energy dependence of the ν_μ charged-current cross-section is included in the calculation but is most relevant for the small off-diagonal elements. As with the F/N ratio, the Far Detector spectrum predicted by the matrix is relatively insensitive to uncertainties in the hadron production calculation, neutrino cross-sections and detector effects.

The matrix can only be employed after accounting for detector acceptance and inefficiencies. The Near Detector measurement is first corrected to remove non- ν_μ charged-current contributions. A second correction, derived from a sample of fully simulated events in the Near Detector, is then applied to deconvolve the detector efficiency and energy resolution. Both corrections are derived from the simulation described in Sec. III E. The resulting distribution estimates the true ν_μ charged-current energy spectrum at the Near Detector. This distribution, organized in energy bins, is treated as an m -dimensional column vector $\{N_i\}$ and multiplied by the $m \times m$ dimensional matrix $\{B_{ij}\}$ to estimate the true ν_μ charged-current energy at the Far Detector

$$F_i = \sum_{j=0}^m B_{ij} N_j \quad \text{“The Beam Matrix Method”} \quad . \quad (9)$$

The effects of neutrino oscillations may be applied to the derived Far Detector spectrum after which detector energy resolution, efficiency, and non- ν_μ charged-current backgrounds are added back in. The resulting spectrum may be directly compared to the Far Detector data.

The direct extrapolation methods described here reduce, but do not entirely eliminate, the effect that uncertainties in hadron production, neutrino cross-sections and detector acceptance have on the prediction of the neutrino energy spectrum at the Far Detector. We will discuss these uncertainties in greater detail in Sec VI C.

2. Indirect Extrapolation Methods

In these methods, a fit is performed to observed Near Detector distributions using a Monte Carlo dataset with parametrized uncertainties due to beam modeling, neutrino cross-sections and neutrino energy scales. The result of




FIG. 31: The joint distribution of neutrino energies observed in the Near and Far Detectors. The contents of each cell represent the mean number of ν_μ events expected in the Far Detector for one event in the Near Detector. This distribution may be treated as a matrix, as in Eq. 9, to relate the energy spectra measured in the Near Detector to those in the Far Detector.

this fit is a tuned Monte Carlo calculation which can then be applied to the Far Detector in order to obtain the predicted neutrino spectrum. These fits have the advantage that they can, in principle, separate systematic effects according to their effect on the fitted distributions with the disadvantage that the fitting functions may not adequately describe the data or (due to degeneracies in the fitted parameters) correctly describe the underlying physics.

Two indirect extrapolation methods have been developed: “NDFit” and “2DFit”. In both techniques a fit is performed to the measured Near Detector energy spectra in order to improve agreement between the observed distributions and those predicted from Monte Carlo simulation. Both methods use data collected in the six beam configurations of Tab. I with approximately equal numbers of events in each configuration. In total, approximately 2.5×10^5 Near Detector data events are used. The fitting techniques are permitted to re-extract the neutrino energy scale and normalization parameters originally determined when tuning the flux calculation (see Tab. IV). The two techniques employ very similar fitting procedures. In the case of the NDFit method, the fit is performed to one-dimensional reconstructed neutrino energy histograms, whereas in the case of the 2DFit, the reconstructed neutrino energy distributions are sub-divided in bins of inelasticity, $y = 1 - E_\mu/E_\nu$.

The NDFit method varies the parameters $\alpha = \{\alpha_j\}$ to minimize

$$\chi^2 = \sum_k^{E_\nu\text{-bins}} \frac{(N_k - n_k(\alpha))^2}{\sigma_k^2 + S_k^2} + \sum_{j=1}^{\text{syst}} \left(\frac{\Delta\alpha_j}{\sigma_j^\alpha} \right)^2 \quad \text{“The NDFit Method”} \quad , \quad (10)$$

where N_k and n_k are the numbers of data and simulated events in bin k of the reconstructed neutrino energy distribution. There are 38 energy bins for each of the six beam configurations (see Tab. I and Fig. 26). The $\{\Delta\alpha_j\}$ are the deviations from nominal of the systematic parameters α , with associated uncertainties $\{\sigma_j^\alpha\}$. These are described below. The second term in Eq. 10 constrains the fit by increasing χ^2 as the parameters are varied away from their nominal values. The error on the number of predicted events is the sum in quadrature of the statistical $\{\sigma_k\}$ and systematic $\{S_k\}$ uncertainties in the neutrino flux, where the latter are derived from the calculation described in Sec. V.

The five systematic parameters $\alpha = \{\alpha_1 \dots \alpha_5\}$ describe uncertainties in the shape and normalization of the ν_μ charged-current cross-section and the muon and hadron energy scales of the detector. These parameters are: α_1 , charged-current axial vector masses for quasi-elastic and resonance production processes, varied coherently; α_2 , cross-section scale factor for the non-resonant 1π and 2π production cross-section at invariant masses $W < 1.7 \text{ GeV}/c^2$; α_3 , the absolute hadronic energy scale; α_4 , the absolute muon energy scale; α_5 , the overall normalization. The magnitude of these uncertainties is shown in Tab. V. The hadronic energy scale uncertainty is dominated by final state interactions (see Sec. III E and III F). The muon energy scale was estimated from a comparison of range and curvature measurements for stopping muon tracks observed in the Near Detector. The normalization uncertainty accounts for the remaining uncertainties left after the improved flux calculation of Sec. V.

The result of the NDfit is shown in Fig. 32. Here the ratio of observed over simulated reconstructed neutrino energy spectra are shown before and after the fit for three of the six beam configurations. The large excursions in the ratio formed with the nominal simulation are substantially reduced by the improved hadron production calculation of Sec. V. Further improvements are observed after minimizing Eq. 10, especially in the 0-5 GeV region of the LE10/185kA configuration, and around the focusing peak of the LE100/200kA and LE250/200kA spectra. Initial excursions in the data/MC ratio of $\sim 30\%$ have been reduced to $\sim 10\%$ or less in all six spectra. Remaining deviations are generally within the uncertainties of the hadron production calculation.

Table V shows the best-fit values of the five systematic parameters used in the NDfit. All five parameters lie within one standard deviation of their nominal values. It is important to note that the results do not constitute a direct measurement of these quantities, as strong correlations exist between the fit variables. The value of χ^2 at the best-fit point is 186.0 for 228 degrees of freedom, of which the penalty terms contribute four units of χ^2 for five parameters.

	Systematic parameter	Uncertainty	Best-fit
α_1	CC+RES axial-vector mass	5%	+5%
α_2	non-resonant 1 and 2π production	20%	+12%
α_3	Hadronic energy scale	10%	-4%
α_4	Muon energy scale	2%	-3%
α_5	Normalization	5%	-2.5%

TABLE V: Systematic parameter values returned from the NDFit (Eq. 10).

The 2DFit method treats the Near Detector data as a function of both the reconstructed neutrino energy E_ν and inelasticity y in order to allow degeneracies between muon and hadron energy scale parameters to be broken. In the fit, the joint distribution of reconstructed neutrino energy and y is divided into a two dimensional grid with 40 energy bins spaced between 0 GeV and 30 GeV and eight reconstructed y bins between zero and one. Data from each of the six beam configurations is tabulated separately, yielding a total of 1920 analysis bins which are employed simultaneously in the fit. The fit attempts to minimize

$$\chi^2 = \sum_k^{E_\nu\text{-bins}} \sum_l^{y\text{-bins}} \frac{(N_{k,l} - n_{k,l}(\beta))^2}{\sigma_{k,l}^2} + \sum_{j=1}^{\text{syst}} \left(\frac{\Delta\beta_j}{\sigma_j^\beta} \right)^2 \quad \text{“The 2DFit Method”} \quad , \quad (11)$$

where the variables $N_{k,l}$ and $n_{k,l}$ refer to the contents of the E_ν, y bins for the selected ν_μ charged-current data and Monte Carlo samples respectively. Like the α_j in Eq. 10, the β_j are parameters which account for the systematic uncertainties in the Monte Carlo simulation and are varied in the fitting procedure to improve agreement with the data.




FIG. 32: Ratio of Near Detector data to Monte Carlo simulation as a function of reconstructed neutrino energy for three beam spectra before and after the NDfit. The dotted gray histograms show the ratio relative to the nominal simulation, the dashed black lines show the ratio after applying the improved hadron production calculation (Sec. V), and the black points with error bars show the ratio after the best-fit NDfit parameters are also applied to the simulation. The errors on the data points represent statistical uncertainties and the shaded regions represent the residual systematic uncertainty from the beam-re-weighting procedure.

The $n_{k,l}$ are implicit functions of the parameters β which account for the following systematic uncertainties: β_1 , the normalization of the neutral-current background; β_2 , an energy dependent variation in the event rate dictated by the uncertainty in the tuned hadron production calculation; β_3 , the absolute hadronic energy scale; β_4 the absolute muon energy scale; In addition, the true neutrino energy versus y distribution for Monte Carlo events is divided into a two dimensional grid composed of eight neutrino energy bins (0-2, 2-4, 4-6, 6-8, 8-10, 10-15, 15-20, > 20 GeV) and five y bins (0.0-0.2, 0.2-0.4, 0.4-0.6, 0.6-1.0). This yields an additional 40 fit parameters, $\beta_5 \dots \beta_{44}$, which allow the normalization of the Monte Carlo to vary independently in each bin with an uncertainty of 20%. This uncertainty is comparable to the *a priori* uncertainties in the neutrino flux and charged-current cross-sections. Table VI shows the magnitude of the systematic parameters. In total, 45 fit parameters are varied to minimize Eq. 11.

Systematic parameter	Uncertainty	Best-fit
β_1	Neutral-current background	50%
β_2	Flux uncertainty	energy dependent
β_3	Hadronic energy scale	10%
β_4	Muon energy scale	2%
$\beta_5 \dots \beta_{44}$	(E_ν, y) normalization factors	20%
		-37%
		0.88σ
		+2.5%
		-4.8%
		0.64-1.61

TABLE VI: Systematic parameter values returned from the 2DFit (Eq. 11). The uncertainty on the flux parametrized by β_2 is energy dependent, approximately $\pm 9\%$ ($\pm 6\%$) for $E_\nu \lesssim 6$ GeV ($E_\nu \gtrsim 6$ GeV). The fit parameter is expressed in units σ of the uncertainty. The parameters $\beta_5 \dots \beta_{44}$ take values between 0.64 and 1.61 with a mean of 1.02 and an r.m.s. of 0.21, approximately equal to the *a priori* uncertainty.

Figure 33 shows the neutrino energy distribution from the LE10 185 kA configuration split into four ranges of y . The value of χ^2 at the best-fit point is 2606.3 for 1919 degrees of freedom, compared to a χ^2 of 5432.1 before the

Extrapolation Method	Predicted Number of Far Detector Events
Beam Matrix	336.0 ± 14.4
Far/Near Ratio	330.0 ± 14.2
NDFit	332.7 ± 14.3
2DFit	323.9 ± 13.9

TABLE VII: Predicted numbers of Far Detector events for the four extrapolation methods for an exposure 1.27×10^{20} POT. The systematic errors on the predictions, which are dominated by the 4% relative normalization uncertainty, are also shown.

fit. The best-fit values of the 40 (E_ν, y) normalization factors range from 0.64 to 1.61, with a mean of 1.02 and an r.m.s. of 0.21 (approximately equal to the *a priori* uncertainty). The best-fit values for the other parameters are -37% for neutral-current normalization and -4.8% and +2.5% for muon and shower energy scales, respectively. The overall normalization was reduced according to a 0.88σ shift in the tuned hadron production calculation.

B. Predicted Far Detector spectrum

Figure 34a shows a comparison of the predicted Far Detector unoscillated visible energy spectra obtained from the four extrapolation methods, for an exposure of 1.27×10^{20} protons on target. The ratios of the three other spectra relative to the Beam Matrix prediction are shown in Fig. 34b. The agreement between the predicted spectra is better than $\pm 5\%$ in the region 1-15 GeV, where the oscillation signal is expected to lie. The spread on the predicted spectra is significantly smaller than the statistical error when the data are binned in 1 GeV wide energy bins. Table VII shows the number of Far Detector events predicted between 0-30 GeV for each extrapolation method. The systematic errors on the predicted rate, which are described below, are dominated by the 4% relative normalization error.

C. Sensitivity to systematic errors for the different extrapolation methods

In this section we describe the various sources of systematic uncertainty and their impact on the neutrino spectrum predicted at the Far Detector. Each source was considered separately and Monte Carlo datasets were generated with $\pm 1\sigma$ shifts of the systematic parameter applied along with oscillations according to $\Delta m^2 = 2.72 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta = 1$, chosen to agree with our best fit to the data. These “mock datasets” were produced for both Near and Far Detectors and were analyzed in the same way as the actual data. In particular the tuning procedure of Sec. V was run on each dataset and the extrapolation methods were able to use the Near Detector mock data to predict the spectra at the Far Detector. These predictions were then used to fit the Far Detector mock data for neutrino oscillations, allowing us to examine the way in which systematic uncertainties affect our measurements of Δm^2 and $\sin^2 2\theta$. The largest sources of uncertainty were identified with this procedure and then accounted for when fitting the Far Detector data for neutrino oscillations.

1. Systematic error sources

Figure 35 illustrates the effect that four representative sources of uncertainty have on the Far Detector neutrino energy spectrum, and the way in which the extrapolation methods are able to account for them with Near Detector measurements. Black points show the ratio of reconstructed energy spectra at the Far Detector, where the denominator corresponds to the nominal simulation and the numerator was generated with a particular systematic uncertainty source shifted with respect to the default value. Deviations from unity indicate the way in which the systematic uncertainties modify the neutrino energy spectrum. The size of the deviations does not necessarily indicate the potential effect on Δm^2 and $\sin^2 2\theta$ because the extrapolation methods use similarly shifted Near Detector Monte Carlo samples as “mock data”. Their prediction of the ratio, shown as lines, would follow the black points if the systematic shift could be completely corrected with Near Detector data. In that case one would extract the same oscillation parameters as were put into the study and the particular source of uncertainty would not affect the results. We describe the sources of uncertainty below.

Charged-current cross-sections: We assign a $\pm 10\%$ uncertainty in the axial vector masses for quasi-elastic and resonance production processes. Fig. 35a shows the effect on the FD spectrum when a +10% shift in those quantities






FIG. 33: Near Detector reconstructed neutrino energy spectra collected in the LE10/185kA beam configuration and for different reconstructed y regions. Open circles show the data. The Monte Carlo simulation is shown before and after the tuning procedure performed by the 2DFit technique (Eq. 11). To simplify the presentation we have grouped the eight reconstructed y bins used in the fit into four bins here.



(a)



(b)

FIG. 34: Neutrino energy spectra at the Far Detector in the absence of neutrino oscillations as predicted by the four extrapolation methods.

is applied in the simulation. The predicted spectra from the four extrapolation methods successfully reproduce the shift. This is expected, since cross-section changes are common to both Near and Far detector events, the Near to Far extrapolation methods should provide a significant cancellation of such uncertainties. We have also conducted a similar study to verify that our analyses are able to correct for the $\pm 20\%$ uncertainty in the magnitude of the non-resonant 1π and 2π production cross-sections for invariant masses $W < 1.7 \text{ GeV}/c^2$. See Sec. III E for further details.

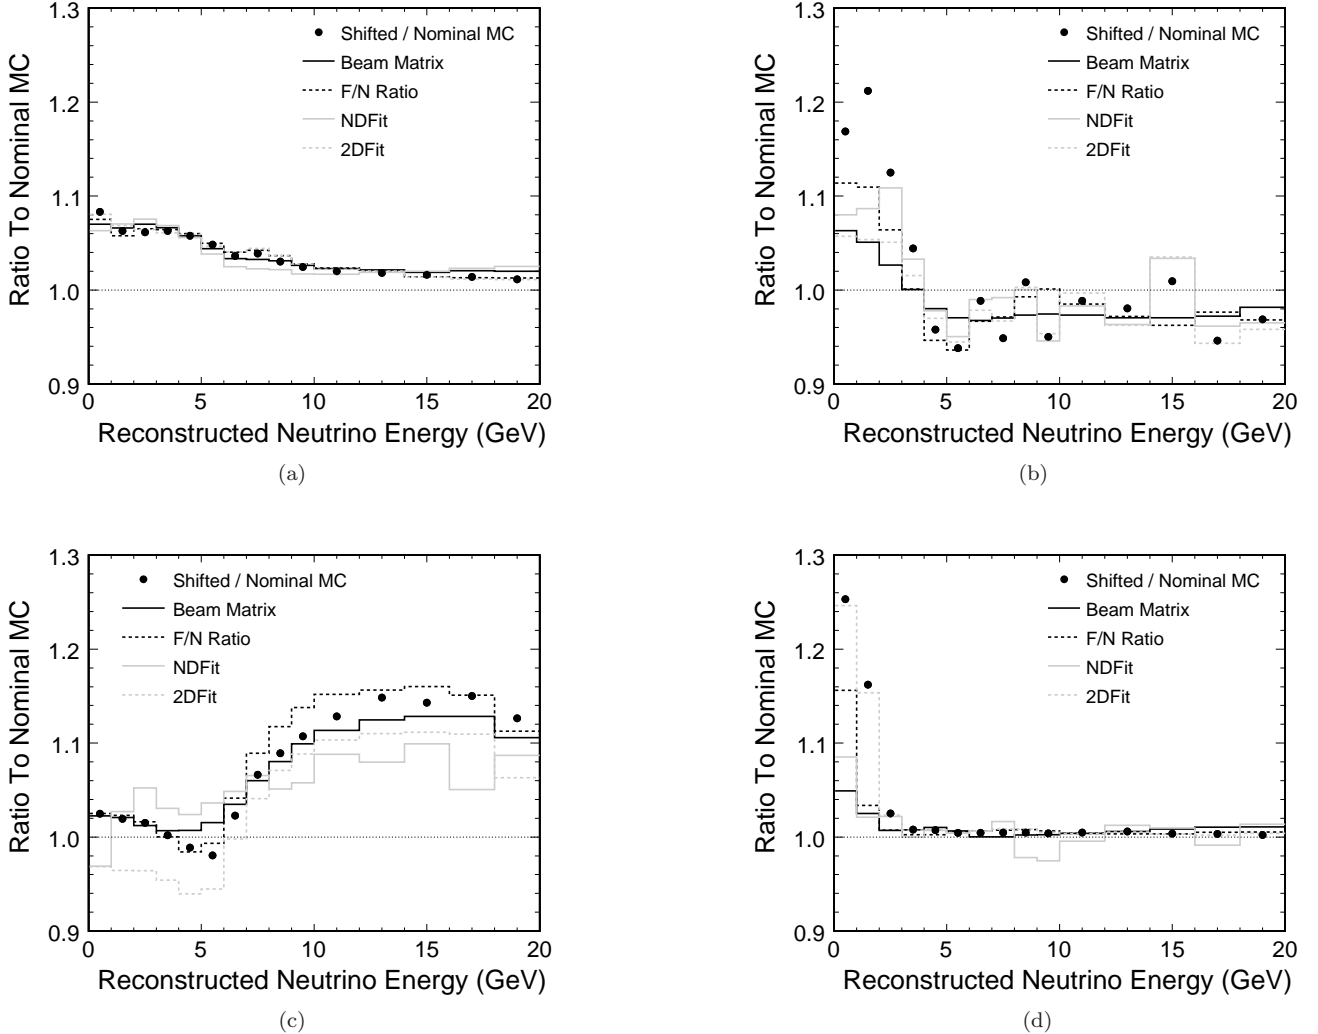


FIG. 35: Effect of systematic uncertainties on the predicted Far Detector energy spectrum. The following systematic shifts are shown: (a) $+10\%$ increase in ν_μ charged-current quasi-elastic and resonance production cross-sections; (b) -10% shift in reconstructed shower energy scale; (c) $+50\%$ uncertainty on the beam tuning correction; (d) $+50\%$ increase in neutral-current background rate. Further details are described in the text.

Final-state interactions: As described in Sec. III E we estimate that uncertainties in final-state interactions cause a $\pm 10\%$ uncertainty on the shower energy scale. Figure 35b shows the effect that a -10% shift in the energy scale has on the Far Detector neutrino spectrum. The distortion of the spectrum below 10 GeV is significant. Because the same shift is present in both Near and Far Detectors, the predicted spectra from the four methods provide an improved description of the distorted spectrum and thus reduce the effect of this systematic error. However, the neutrino energy and Δm^2 are coupled in Eq. 1 and a change in the energy scale may be incorrectly interpreted as a shift in Δm^2 . It is therefore necessary to explicitly account for energy scale uncertainties when extracting the oscillation parameters from our data.

Beam-related uncertainties: The beam tuning procedure described in Sec. V applies corrections to the (p_z, p_T) distribution in the Monte Carlo that result in changes in the visible energy spectrum of approximately 10% around the peak in the LE10/185kA beam configuration and about 30% in the high energy tail. We have conservatively assigned a 50% error on the magnitude of this correction as an estimate of the uncertainty on the beam tuning. The effect of this uncertainty on the FD spectrum is shown in Fig. 35c. There is a significant shift in the high-energy tail of the spectrum ($E_\nu > 5 \text{ GeV}$) where the beam tuning correction is largest. The plot shows that the direct extrapolation methods are in general more robust than the fitting methods to this type of uncertainty, as they are much less dependent on the requirement that the tuned Monte Carlo closely match the data in order to provide an accurate prediction of the FD spectrum.

Normalization: We estimate an uncertainty of $\pm 4\%$ in the relative normalization of the energy spectra measured in the Near and Far Detectors. The estimate is composed from a 2% uncertainty in the fiducial mass of both detectors, a 3% uncertainty in the relative Near and Far Detector reconstruction efficiencies, estimated from a visual scan of Near and Far Detector data and Monte Carlo events, and a 1% uncertainty in the live time.

Neutral-current background: As described in Sec. IV B 1, we estimate an uncertainty of $\pm 50\%$ on the rate of neutral-current events misidentified as ν_μ charged-current. Figure 35d shows that the effect of this uncertainty is largest in the lowest reconstructed energy bins ($E_\nu < 2 \text{ GeV}$) and is only partially corrected by the F/N, Beam Matrix and NDFit extrapolation methods. The 2DFit method includes the neutral-current rate as a fit parameter and recovers the correct prediction in this special case. In the general case however, where several systematic effects will be present, no method will be able to completely correct for this source of uncertainty.

Shower energy scale calibration: Based on a comparison of test beam measurements in the Calibration Detector and the Monte Carlo simulation (see Sec. III F) we estimate a 6% uncertainty in the response to single pions and protons. The uncertainty in the response to electrons and muons is negligible. As described in Sec. III F, stopping muons were used to cross-calibrate the two detectors with a relative uncertainty of $\pm 2\%$.

Muon energy scale: A 2% muon energy scale uncertainty is assumed based on studies of fully contained muon tracks reconstructed in the Near Detector. The difference between the momentum obtained from the track range and the momentum obtained from a fit to the curvature of the track due to the magnetic field of the detector was examined for both real and simulated events, as a function of the range momentum. The deviation of this quantity between data and simulation was approximately 2%, which was taken as an estimate of the uncertainty on the magnetic field calibration in the Near Detector, and hence the error on the relative muon energy scale between Near and Far Detectors. Our estimate utilizes the muon range/energy relation which has an uncertainty of approximately 2% based on a material assay of the detectors and a comparison of our muon stopping power tables with those in [32].

As discussed in Sec. III E, our magnetic field was recalibrated after we performed this analysis. This predominantly affects the reconstruction of muons with a momentum larger than $\sim 7 \text{ GeV}/c$ which arise from neutrinos outside the energy region in which we observe oscillations.

CC selection efficiency: We varied the requirement on the event selection parameter S (see Eq. 5) for mock data sets by ± 0.02 while holding the cut applied in the nominal Monte Carlo dataset constant. This changes both the number of true charged-current and neutral current events that are classified as ν_μ charged-current. The magnitude of this shift was obtained from a comparison of the data and Monte Carlo S distributions in the Near Detector. We used mock datasets including the systematic uncertainties listed above to examine the accuracy with which the different methods are able to predict the Far Detector spectrum and study the $\sin^2 2\theta$ and Δm^2 extracted from our oscillation fits. As described below, this procedure was used to explore the capabilities of the different extrapolation techniques and identify the most important sources of systematic uncertainty.

2. Choice of primary analysis method

The indirect extrapolation methods (NDFit, 2DFit) attempt to adjust the Monte Carlo prediction of the neutrino energy spectrum in the Near Detector to improve agreement with the data and then use the improved Monte Carlo calculation to predict the Far Detector spectrum. These procedures are not able to arrive at a perfect description of the data. In particular, discrepancies of $\pm 5\%$ are present when data collected in the LE10/185kA beam configuration are compared with the Monte Carlo prediction after the best-fit NDFit and 2DFit systematic parameters have been applied. We also notice that somewhat larger discrepancies are still present in the LE100/200kA beam configuration. If the major source of these distortions is caused by uncertainties in the neutrino flux or neutrino cross-sections then the direct (Beam Matrix and the F/N) extrapolation methods simply translate the measurements to the Far Detector and the discrepancies do not affect the oscillation measurement. As a consequence we have found that the predictions made by the direct methods are less sensitive to the absolute agreement between data and the Monte Carlo simulation in the Near Detector than are the those made by the indirect methods.

In light of these issues, we have decided to use one of the direct extrapolation methods to obtain our primary

oscillation result. We chose the Beam Matrix method as our primary extrapolation technique because it had smaller systematic errors than the F/N ratio method when all sources of uncertainty were considered. The results that we will present in Sec. VII therefore use the Beam Matrix method to predict the Far Detector spectrum and perform the oscillation fit. As a cross-check of these results, we present the best fit oscillation parameters and allowed regions obtained from the other three extrapolation methods as well.

3. Systematic Uncertainties in the Oscillation Fit

The three largest contributions to the systematic error for the Beam Matrix method are a) the uncertainty in the relative normalization of the energy spectra measured in the two detectors, b) uncertainties in the absolute hadronic energy scale, and c) uncertainties in the neutral current background rate. The systematic shifts calculated for the other sources of uncertainty are small in comparison. The magnitude of these shifts is summarized in Tab. VIII. As expected from the above discussion, because uncertainties due to beam modeling and cross-sections are common to the two detectors they largely cancel out in the extrapolation. The three largest uncertainties are included as systematic nuisance parameters in the Far Detector oscillation fit. By fitting for these systematic parameters simultaneously using both Near and Far Detector data, the effect of these uncertainties is substantially reduced, due to significant cancellations of these errors between the two detectors.

VII. OSCILLATION ANALYSIS

With an exposure of 1.27×10^{20} POT, a total of 215 beam-coincident events with reconstructed energies below 30 GeV are selected as ν_μ charged-current in the Far Detector. Assuming no oscillations, the predicted number of Far Detector events in the same energy range for this exposure is 336 ± 14 . The error quoted here is dominated by the 4% systematic error on the overall normalization. The deficit corresponds to a significance of 5.2 standard deviations, where both statistical and systematic errors on the total rate are taken into account. In this section, we describe an oscillation analysis of the observed Far Detector reconstructed energy spectrum. We present results obtained using the extrapolated Far Detector spectrum from the Beam Matrix method, and compare these with results from the other three extrapolation methods.

Figure 36a shows the reconstructed energy distribution of the selected Far Detector events. A fit to these data is performed to extract the mixing parameters Δm^2 and $\sin^2 2\theta$, within the context of two-flavor $\nu_\mu \leftrightarrow \nu_\tau$ oscillations (Eq. 1). We minimize the following statistic, $\chi^2 = -2 \ln \lambda$, where λ is the likelihood ratio [60]:

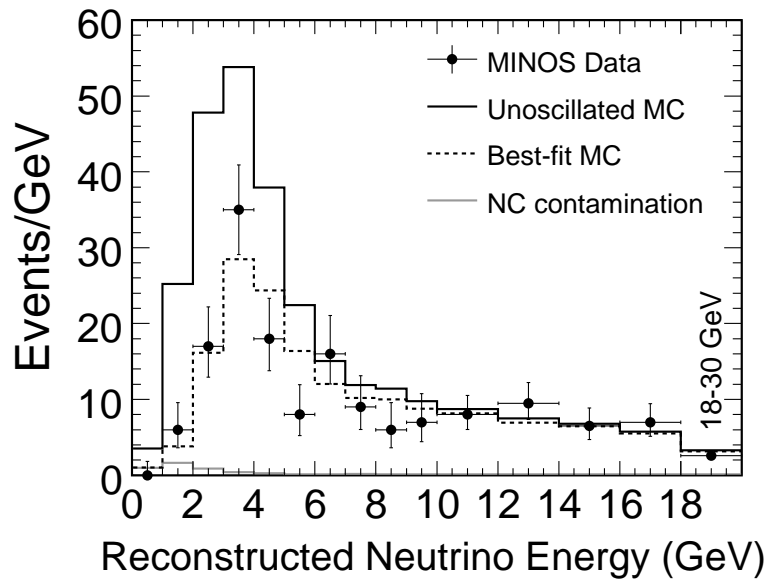
$$\chi^2(\Delta m^2, \sin^2 2\theta, \alpha_j) = \sum_{k=1}^{15} 2(N_k^{exp} - N_k^{obs}) + 2N_k^{obs} \ln(N_k^{obs}/N_k^{exp}) + \sum_{j=1}^3 \left(\frac{\Delta \alpha_j}{\sigma_{\alpha_j}} \right)^2, \quad (12)$$

where N_k^{obs} and N_k^{exp} are the numbers of observed and expected events in bin k of the reconstructed energy distribution, the α_j are fitted systematic parameters, with associated errors σ_{α_j} .

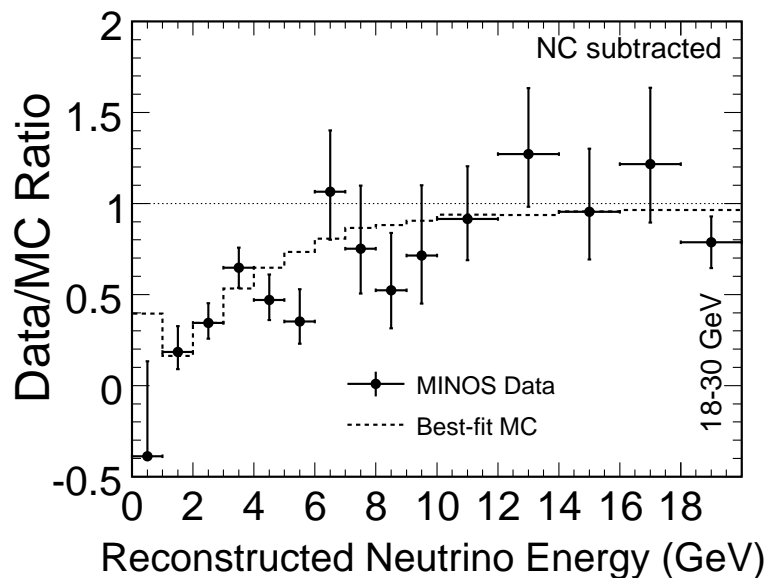
The three leading systematic uncertainties identified in Sec. VI C are included in the fit as nuisance parameters, with Gaussian distributed errors. These parameters are the relative normalization between the Far and Near Detectors, with a 4% uncertainty; the absolute hadronic energy scale with a 11% uncertainty, and a 50% uncertainty in the neutral-current background rate for the selected sample. Since the energy scale and neutral-current background uncertainties are common to the two detectors, their effect significantly cancels in the extrapolation of the energy spectrum from the Near to the Far Detector. In order to account for this in the fit, these parameters are varied simultaneously for both Near and Far Detectors, and the Monte Carlo reconstructed energy distributions are modified accordingly. A new Far Detector predicted spectrum is then obtained for every value of these systematic parameters as they are varied in the oscillation fit.

Figure 36a shows the reconstructed neutrino energy spectrum measured at the Far Detector, the unoscillated Far Detector predicted spectrum obtained from the Beam Matrix method, along with the predicted spectrum weighted by the best-fit oscillation parameters. Figure 36b shows the ratio of data to the unoscillated Monte Carlo prediction as a function of reconstructed energy, and the predicted ratio for the best-fit oscillation parameters. The contamination from misidentified neutral-current events, which is shown by the gray histogram in Fig. 36a, is subtracted from both the predicted and observed distributions in Fig. 36b. The shape of the data distribution is well modeled by the oscillation hypothesis.

The allowed regions at 68, 90, 99% C.L. in the $\Delta m^2, \sin^2 2\theta$ plane from a fit to the 215 Far Detector selected data events using the Far Detector predicted spectrum are shown in Fig. 37. Here the confidence level intervals are obtained



(a)



(b)

FIG. 36: (a) The reconstructed energy spectra of selected Far Detector events with the Far Detector unoscillated prediction (solid histogram) and best-fit oscillated spectrum (dashed histogram) overlaid. The predicted neutral-current background in the selected sample is shown in gray. The right-most bin in this distribution contains all events between 18 and 30 GeV. The asymmetric error bars on the data points represent the 68% C.L. Poisson errors on the numbers of observed events. (b) The ratio of the observed spectrum to the unoscillated Far Detector prediction, where the expected neutral-current background has been subtracted.

using the Gaussian approximation ($\Delta\chi^2 = 2.3, 4.6, 9.2$). These confidence-level intervals were found to be in good agreement with those obtained from a study using the unified approach of Feldman and Cousins [61]. The best-fit parameters are $\Delta m^2 = 2.74 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta = 1$, where the fit has been constrained to the region $\sin^2 2\theta \leq 1$. The allowed ranges of these parameters at 68% C.L. and for 1 d.o.f. are $2.54 \times 10^{-3} < \Delta m^2 < 3.18 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta > 0.87$. The χ^2 at the best-fit point is 20.3 for 13 degrees of freedom, which corresponds to a χ^2 probability $P(\chi^2, \text{d.o.f.}) = 8.9\%$. It has been verified using Monte Carlo experiments that the probability of 8.9% is valid for the relatively low statistics of the data sample. The χ^2 for the null oscillation hypothesis is 104 for 15 degrees of freedom.

The contribution of ν_τ charged-current events from $\nu_\mu \leftrightarrow \nu_\tau$ oscillations is included in the fit. At the best-fit point, the expected number of ν_τ charged-current events in our sample is 0.78 events. The expected background from oscillated ν_e charged-current events at our best-fit point is < 0.3 events, for $\nu_\mu \leftrightarrow \nu_e$ mixing with $\sin^2 2\theta_{13} < 0.2$ [62]. This background, and the 0.38 events expected from rock muon interactions, are not included in the fit.




FIG. 37: Allowed regions at 68,90,99% C.L. in the $\Delta m^2, \sin^2 2\theta$ plane from a fit to the Far Detector reconstructed energy spectrum using the Beam Matrix extrapolation method. The best-fit point, which occurs at $\Delta m^2 = 2.74 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta = 1$, is represented by the star.

A fit to the Far Detector energy spectrum where the physical boundary constraint ($\sin^2 2\theta \leq 1$) is removed yields best-fit parameters that are very slightly in the unphysical region: $\Delta m^2 = 2.72 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta = 1.01$, with a best-fit $\chi^2/\text{d.o.f.} = 20.3/13$. The one-dimensional projection of the χ^2 surface for $\sin^2 2\theta$, where the value of χ^2 has been minimized at each point with respect to Δm^2 , is shown in Fig. 38.

Figure 39 shows the effect of systematic errors on the measurement of the oscillation parameters. The figure shows the 90% C.L. allowed regions obtained from fits to the Far Detector data assuming statistical errors only, and




FIG. 38: One-dimensional projection of the $\Delta\chi^2$ surface for $\sin^2 2\theta$ using the Beam Matrix extrapolation method, for a fit where the physical boundary at $\sin^2 2\theta = 1$ has been removed.

statistical and systematic errors combined. The 90% C.L. allowed region increases in size by approximately 10% in both Δm^2 and $\sin^2 2\theta$ when systematic errors are also taken into account. This indicates that, for this exposure, the measurement errors on the oscillation parameters are limited by statistical uncertainties. Table VIII shows the systematic shifts on the best-fit point for the various sources of systematic error considered in Section VIC. The shifts due to beam and cross-section uncertainties are negligibly small. As an additional check we have repeated the analysis without the improved flux calculation of Sec. V and find that Δm^2 changes by $2 \times 10^{-5} \text{ eV}^2/c^4$, consistent with the uncertainty quoted in Tab. VIII. The shifts due to the three largest systematic errors combined (a)–(c) in Tab. VIII) are approximately three times smaller than the statistical errors on Δm^2 and $\sin^2 2\theta$. The magnetic field recalibration discussed in Sec. III E predominantly affects the portion of the neutrino spectrum where oscillations do not occur and may be closely approximated by scaling the strength of the field used in analysis of the data. This results in shifts in the best fit Δm^2 and $\sin^2 2\theta$ that are small in comparison with the major systematic errors in Tab. VIII.

Uncertainty		$ \Delta m^2 $ ($10^{-3} \text{ eV}^2/c^4$)	$\sin^2 2\theta$
(a) Normalization	($\pm 4\%$)	0.050	0.005
(b) Abs. hadronic E scale	($\pm 11\%$)	0.057	0.048
(c) NC contamination	($\pm 50\%$)	0.090	0.050
(d) Beam uncertainties		0.015	< 0.005
(e) Cross sections		0.011	0.005
All other systematics		0.041	0.013
Statistical Error		0.35	0.13

TABLE VIII: Systematic shifts on the measurement of $|\Delta m^2|$ and $\sin^2 2\theta$ for various sources of systematic error, calculated using Monte Carlo generated with the best-fit oscillation parameters obtained from the Far Detector data and the Beam Matrix extrapolation method. The values quoted are the average shifts for ± 1 standard deviation variations in each of the systematic parameters. The last row of the table shows the expected statistical uncertainty on the measurement of $|\Delta m^2|$ and $\sin^2 2\theta$ for an exposure of 1.27×10^{20} POT. The shifts on $|\Delta m^2|$ and $\sin^2 2\theta$ are treated as uncorrelated, and correlations between the various systematic effects are not taken into account.

Figure 40 shows the 90% C.L. allowed regions obtained from fits to the 215 Far Detector data events using the Far Detector predictions obtained from the four extrapolation methods described in the previous section. Both the allowed regions, and the best-fit parameters, which are shown in Tab. IX, are in very good agreement between the methods. The spread in the allowed regions is small relative to the size of the regions — this spread is due to the small differences in the predicted Far Detector spectra shown in the previous section.




FIG. 39: Comparison of the 90% C.L. regions obtained by considering both statistical and systematic errors (black contour) and statistical errors only (gray contour). The best-fit point for the fit including statistical and systematic errors is shown by the star.

The relative sensitivity of the oscillation fit to shape and rate information is illustrated in Fig. 41. This shows the 90% C.L. regions that are obtained from fits to the total rate of Far Detector events, where no spectral information is used and fits to the shape of the spectrum only where data and MC distributions are normalized to the same number of events. The shape of the spectrum plays the most important role in defining the size of the allowed region; the rate information alone does not provide an upper bound on the value of Δm^2 .

Figure 42 shows a comparison of the 90% C.L. allowed region from MINOS with those previously reported from the K2K long-baseline [8] and the Super-Kamiokande atmospheric neutrino [63] oscillation analyses. Note that the MINOS results are for ν_μ charged-current events, whereas the Super-Kamiokande results are for a combined $\nu_\mu + \bar{\nu}_\mu$ dataset. The allowed regions are in good agreement with each other.

VIII. SUMMARY

In this paper we report on an observation of ν_μ disappearance by the MINOS long-baseline neutrino experiment using an accelerator based neutrino beam provided by the Fermilab Main Injector. The data were collected over a nine month period from May 2005 to February 2006, corresponding to a total of 1.27×10^{20} protons on target. The experiment uses two detectors separated by 734 km. The prediction of the unoscillated neutrino flux at the Far

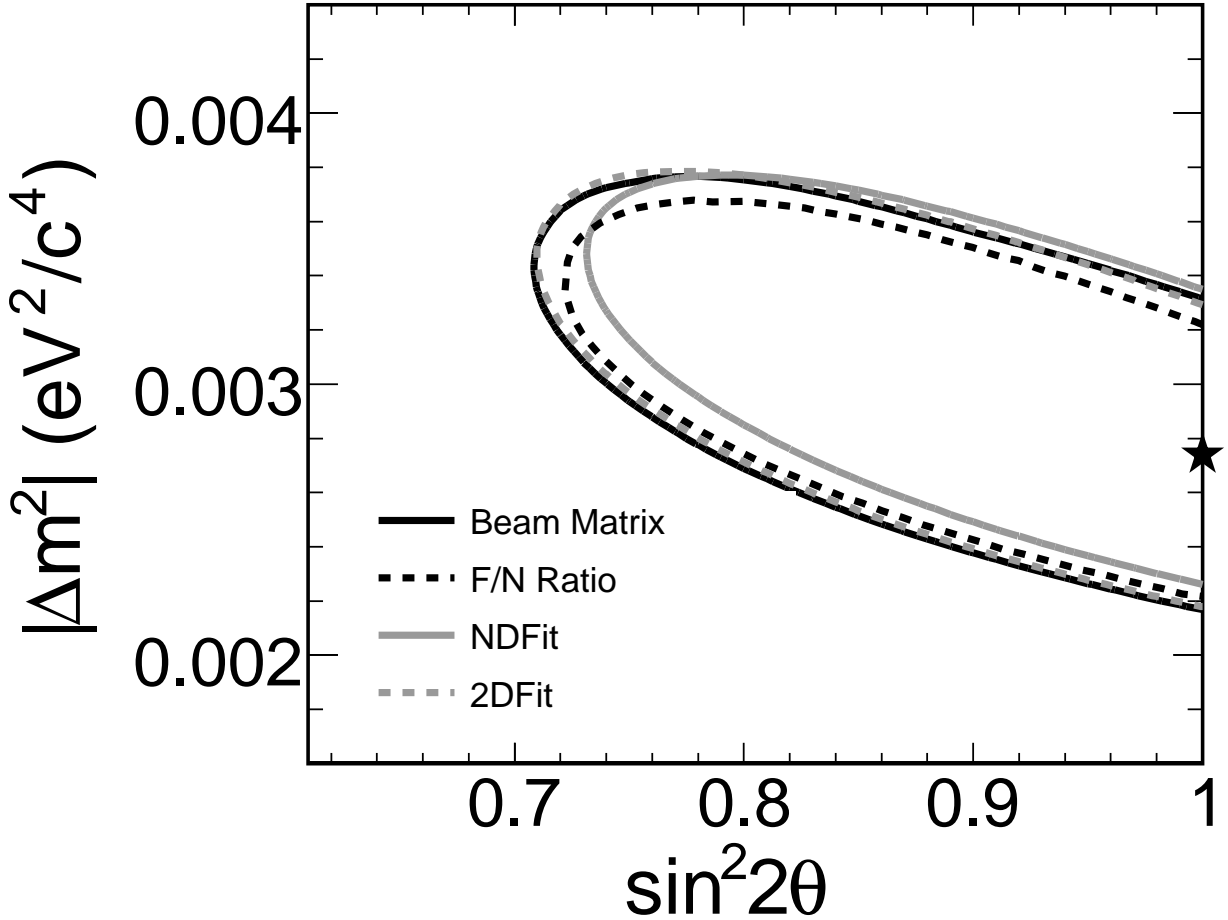


FIG. 40: Comparison of the 90% C.L. regions from oscillation fits using the four extrapolation methods. The best-fit point from the Beam Matrix method is shown by the star. All of the contours include the three leading systematic errors described in the text.

Detector site was obtained from the observed neutrino spectrum at the Near Detector location. This two-detector approach provides significant cancellation of systematic errors due to uncertainties in beam modeling and neutrino cross-sections. In addition, we used Near Detector data taken with several beam configurations in order to constrain some of these uncertainties.

We have used four techniques, each with different sensitivities to systematic uncertainties, in order to obtain the unoscillated Far Detector flux from Near Detector data. All four methods predict very similar Far Detector spectra. The total number of ν_μ charged-current events observed in the Far Detector is 215, compared to the expectation of 336 ± 14 for no oscillations. The deficit in the number of observed events shows a strong energy dependence, consistent with neutrino oscillations.

A fit to the observed energy spectrum, assuming two-flavor $\nu_\mu \leftrightarrow \nu_\tau$ mixing, yields best-fit parameters

$$\Delta m^2 = 2.74 \times 10^{-3} \text{ eV}^2/c^4, \quad \sin^2 2\theta = 1$$

with allowed ranges of $2.54 \times 10^{-3} < \Delta m^2 < 3.18 \times 10^{-3} \text{ eV}^2/c^4$ and $\sin^2 2\theta > 0.87$ (68% C.L., 1 d.o.f.). All four analysis techniques give consistent results. These values are also consistent with those from existing experiments.

The current estimate of our systematic uncertainty is approximately a factor of 2–3 smaller than our statistical error for 1.27×10^{20} protons on target. Continued data taking, together with refinements in the estimation of our systematic errors, will allow us to make significant improvements in our measurements of Δm^2 and $\sin^2 2\theta$ in future

Extrapolation method	Δm^2 (eV^2/c^4)	$\sin^2 2\theta$	$\chi^2/d.o.f$	$\chi^2/d.o.f.$ (no osc.)
Beam Matrix	2.74×10^{-3}	1.0	20.3/13	104/15
Far/Near ratio	2.73×10^{-3}	1.0	52.8/58	132/60
NDFit	2.82×10^{-3}	1.0	20.1/13	96/15
2DFit	2.80×10^{-3}	0.98	34.2/28	107/30

TABLE IX: Best-fit oscillation parameters and χ^2 values returned from fits to the 215 Far Detector data events using four independent extrapolation methods. The right-hand column shows the values of $\chi^2/d.o.f.$ obtained by each method for the null oscillation hypothesis. Fifteen reconstructed energy bins were used by the Beam Matrix and NDFit methods. The F/N ratio used 60 0.5 GeV bins and the 2DFit employed 30 1.0 GeV bins.




FIG. 41: Comparison of the 90% C.L. regions from oscillation fits using shape and rate information only. The best-fit point and 90% C.L. contour from the fit to shape and rate information is also shown.

analyses. Since the initial publication of these results in [1], MINOS has accumulated a total of 3.5×10^{20} POT through July, 2007. Preliminary results from an exposure of 2.5×10^{20} POT recorded in the LE10/185kA beam configuration have been presented in [14] and analysis of the full 3.5×10^{20} POT dataset is in progress.




FIG. 42: Comparison of the 90% C.L. region from MINOS with those previously reported by other experiments.

Acknowledgments

This work was supported by the US DOE; the UK PPARC; the US NSF; the State and University of Minnesota; the University of Athens, Greece and Brazil's FAPESP and CNPq. We are grateful to the Minnesota Department of Natural Resources, the crew of the Soudan Underground Laboratory, and the staff of Fermilab for their contribution to this effort.

-
- [1] D. G. Michael et al. (MINOS), *Phys. Rev. Lett.* **97**, 191801 (2006).
 - [2] B. Aharmim et al. (SNO), *Phys. Rev.* **C72**, 055502 (2005).
 - [3] J. Hosaka et al. (Super-Kamiokande), *Phys. Rev.* **D73**, 112001 (2006).
 - [4] T. Araki et al. (KamLAND), *Phys. Rev. Lett.* **94**, 081801 (2005).
 - [5] Y. Ashie et al. (Super-Kamiokande), *Phys. Rev.* **D71**, 112005 (2005).
 - [6] M. Ambrosio et al. (MACRO), *Eur. Phys. J.* **C36**, 323 (2004).
 - [7] W. W. M. Allison et al. (Soudan-2), *Phys. Rev.* **D72**, 052005 (2005).
 - [8] M. H. Ahn et al. (K2K), *Phys. Rev.* **D74**, 072003 (2006).
 - [9] Z. Maki, M. Nakagawa, and S. Sakata, *Prog. Theor. Phys.* **28**, 870 (1962).
 - [10] B. Pontecorvo, *Sov. Phys. JETP* **26**, 984 (1968).

- [11] H. Nunokawa, S. J. Parke, and R. Zukanovich Funchal, Phys. Rev. **D72**, 013009 (2005).
- [12] K. Anderson et al., Tech. Rep., Fermi National Accelerator Laboratory (2002).
- [13] S. E. Kopp, Proc. 2005 IEEE Particle Accel. Conference (2005), arXiv:physics/0508001.
- [14] P. Adamson et al. (MINOS) (2007), arXiv:0708.1495 [hep-ex].
- [15] M. Bishai et al. (2006), submitted to Nucl. Instr. and Meth. A.
- [16] M. Kostin et al. (2002), FERMILAB-TM-2353-AD.
- [17] A. G. Abramov et al., Nucl. Instrum. Meth. **A485**, 209 (2002).
- [18] G. I. Budker, in *Proceedings of the International Conference on Accelerators* (Dubna, USSR, 1963), pp. 282–283.
- [19] D. G. Baratov et al., Zh. Tekh. Fiz. **47**, 1014 (1977), translated in Sov. Phys. Tech. Phys. **22(5)**, 605 (1977).
- [20] R. Zwaska et al., Nucl. Instrum. Meth. **A568**, 548 (2006).
- [21] S. Kopp et al. (2005), FERMILAB-CONF-05-092, arXiv:physics/0506194.
- [22] S. Kopp et al., Nucl. Instrum. Meth. **A568**, 503 (2006), Fermilab-Pub-06-007.
- [23] V. Bocean (2006), talk given at the 9th International Workshop on Accelerator Alignment (IWAA 06), Menlo Park, California, 26-29 Sep 2006.
- [24] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft (2005), CERN-2005-010.
- [25] R. Brun et al., *GEANT Detector Description and Simulation Tool*, CERN Program Library Long Writeup W5013 (1994).
- [26] D. G. Michael et al., to be submitted to Nucl. Instr. Meth. A.
- [27] P. Adamson et al., Nucl. Instr. Meth. **A492**, 325 (2002).
- [28] A. Belias et al., IEEE Trans. Nucl. Sci. **51**, 451 (2004).
- [29] J. Oliver, N. Felt, G. Feldman, A. Lebedev, and R. Lee, IEEE Trans. Nucl. Sci. **51**, 2193 (2004).
- [30] T. Cuntiff et al., IEEE Trans. Nucl. Sci. **53**, 1347 (2006).
- [31] P. Ballester, Astronomy and Astrophysics **286**, 1011 (1994).
- [32] D. E. Groom, N. V. Mokhov, and S. I. Striganov, Atom. Data Nucl. Data Tabl. **78**, 183 (2001).
- [33] P. Adamson et al. (MINOS), Phys. Rev. **D76**, 072005 (2007).
- [34] Durham Database Group, *HEPDATA: Reaction Data Database*, URL <http://durpdg.dur.ac.uk/HEPDATA/REAC>.
- [35] H. Gallagher, Nucl. Phys. B Proc. Suppl. **112**, 188 (2002).
- [36] D. Rein and L. M. Sehgal, Ann. Phys. **133**, 79 (1981).
- [37] A. Bodek and U. K. Yang, AIP Conf. Proc. **721**, 358 (2004).
- [38] Z. Koba, H. B. Nielsen, and P. Olesen, Nucl. Phys. **B40**, 317 (1972).
- [39] R. D. Ransome, Nucl. Phys. B Proc. Suppl. **139**, 208 (2005).
- [40] R. Merenyi et al., Phys. Rev. **D45**, 743 (1992).
- [41] M. Kordosky, Nucl. Phys. B Proc. Suppl. **159**, 223 (2006), hep-ex/0602029.
- [42] M. A. Kordosky, Ph.D. thesis, University of Texas (2004).
- [43] P. L. Vahle, Ph.D. thesis, University of Texas (2004).
- [44] J. J. Hartnell, Ph.D. thesis, University of Oxford (2005).
- [45] A. Cabrera, Ph.D. thesis, University of Oxford (2004).
- [46] C. Zeitnitz and T. A. Gabriel, Nucl. Instr. Meth. **A349**, 106 (1994).
- [47] P. Adamson et al., Nucl. Instr. Meth. **A556**, 119 (2006).
- [48] N. V. Mokhov et al. (1998), nucl-th/9812038.
- [49] A. J. Malensek (1981), FERMILAB-FN-0341.
- [50] M. Bonesini, A. Marchionni, F. Pietropaolo, and T. Tabarelli de Fatis, Eur. Phys. J. **C20**, 13 (2001).
- [51] L. A. Ahrens et al., Phys. Rev. **D34**, 75 (1986).
- [52] R. C. Allen et al., Phys. Rev. **D47**, 11 (1993).
- [53] P. Astier et al. (NOMAD), Nucl. Instrum. Meth. **A515**, 800 (2003).
- [54] K. S. McFarland et al. (NuTeV) (1998), hep-ex/9806013.
- [55] G. P. Zeller et al. (NuTeV), Phys. Rev. Lett. **88**, 091802 (2002).
- [56] C. L. Wang, Phys. Rev. **D10**, 3876 (1974).
- [57] G. Cocconi, L. Koester, and D. Perkins (1961), LBL Report No. UCRL 10022.
- [58] M. Szleper and A. Para (2001), hep-ex/0110001.
- [59] S. E. Kopp, Phys. Rep. **439**, 101 (2007).
- [60] W.-M. Yao et al., J. Phys. G **33**, 1 (2006).
- [61] G. J. Feldman and R. D. Cousins, Phys. Rev. **D57**, 3873 (1998).
- [62] M. Apollonio et al., Eur. Phys. J. **C27**, 331 (2003).
- [63] Y. Ashie et al. (Super-Kamiokande), Phys. Rev. Lett. **93**, 101801 (2004).